Forecasting Research

Forecasting Research Division
Technical Report No. 208

Report on observation impact studies on cloud and visibility analyses and forecasts in Nimrod during November 1996

by

W H Hand

January 1997

Meteorological Office
London Road
Bracknell
Berkshire
RG12 2SZ
United Kingdom

Forecasting Research
Technical Report No. 208

Report on observation impact studies on cloud and visibility analyses and forecasts in Nimrod during November 1996

by

W H Hand
January 1997

Forecasting Research
Meteorological Office
London Road
Bracknell
Berkshire
RG12 2SZ
United Kingdom

© Crown Copyright 1997

Permission to quote from this paper should be obtained from the above Meteorological Office division.

Report on observation impact studies on cloud and visibility analyses and forecasts in Nimrod during November 1996

by

W.H.Hand

January 1997

0. Summary

An observation impact study on Nimrod cloud and visibility analyses and forecasts has been conducted. Computer problems interrupted the trial and a problem with the retrieval of automatic cloud data in the Nimrod system was identified. The type of trial was therefore not as planned for the cloud component. However, the study demonstrated that severely degrading the observing network had a very detrimental effect on the analysis and prediction of low cloud. Visibility analyses and forecasts were found to be more sensitive to the distribution of observations rather than distinctions between automatic and manual.

1. Introduction

Nimrod is a completely automatic and operational system for producing analyses and very short range forecasts up to six hours ahead of all weather parameters with the main ones comprising precipitation, cloud and visibility (the latter two still under trial). Analyses are made using surface observations, radar and satellite imagery together with NWP fields from the mesoscale configuration of the Unified Model (UM). Forecasts are prepared by combining outputs from the UM with extrapolation forecasts from the analysis. For cloud and visibility, persistence of the analysis is also used as an input. Nimrod is therefore very sensitive to the quantity and quality of observational data. The aims of the observation impact experiments are to assess the impact of removing, or using selected observations, on the Nimrod analyses and forecasts of cloud and visibility.

1.1 Cloud

Nimrod produces a three dimensional cloud analysis and forecast every hour. The multi-level cloud analysis is carried out on horizontal levels above sea-level on a 15km grid over the Nimrod domain (figure 1).The procedure brings together a forecast first guess (usually a one hour Nimrod forecast, but may be a model field if the forecast is unavailable), satellite data and surface observations with varying levels of detail. First the satellite derived cloud cover is refined using surface reports, and the forecast first guess amounts at each level are corrected to agree with the satellite cloud top height. A first guess profile is then obtained for each observation location

Figure 1. The Nimrod domain.

and modified to be consistent with the observed cloud layer information. The modifications are spread using a recursive filter on horizontal levels. Finally, the analysed total cover is used to adjust the layer cloud amounts. In the forecast, (Golding 1996 and Hand 1995), cloud is advected along the horizontal layers using UM winds but ensuring that precipitating cloud is advected using the same vectors as for the precipitation forecast. After each hourly forecast step, the advection forecast, persistence of the analysis and the UM cloud forecast are blended together. Each forecast contribution has a separate weight, estimated for each level and forecast lead time from the accuracy of the forecast contribution at T+0, (although the model has only one comparison). The accuracy measure is the error variance of cloud cover.

In the assessment procedure, forecasts and analyses are compared with observations using the nearest pixel value to the observation. For cloud distribution, the parameters - greater than 4 oktas total cloud, cloud below 5000 feet, and cloud below 1000 feet are assessed using the following scores:

Hit rate (HR) = Number of correct forecasts of > 4/8 / Number observed of > 4/8

False alarm rate (FAR) = Number of incorrect forecasts of > 4/8
/ Number of forecasts of > 4/8

Critical Success Index (CSI) =
Number of correct forecasts of > 4/8 / Number observed or forecast of > 4/8

For significant (greater than or equal to 3 oktas) cloud base height below 5000 feet, root mean square factor (RMSF) errors are calculated.

RMSF = exp(RMS OF [ln(F/A)])

where F/A are the forecast divided by observed values at the observation positions, ln is the natural logarithm and exp denotes exponentiation.

If either the forecast or observed values are greater than 3000m then they are set to 3000m. This is to ensure that the scores are not biased towards high cloud bases. RMSF defines a fractional error and has proved to be a reliable measure of forecast performance in Nimrod. The percentage of occasions where the forecast height is within either 30% or 30m of the observed height is also noted.

Although operationally, Nimrod forecasts are also compared with verifying analyses, this has not been done in the observation impact study (OIS) because of the difficulty in defining a "correct" analysis without biasing results. An alternative would have been to do a cross-verification with each run being compared with the other run's analyses, however, there would not have been enough time to do this on an hourly cycle but it may be worth considering in future impact studies.

1.2 Visibility

Nimrod seeks to analyse and predict a range of visibilities from greater than 10km to dense fog less than 100m. A complete description of the analysis and forecasting system can be found in Wright and Thomas, 1996, but is summarised here. The analysis and forecast are performed using temperature and moisture variables, namely liquid water temperature $T_L = T - Lq_L / C_p$, and the total water $q_T = q + q_L$ where q is the water vapour mixing ratio, q_L is the liquid water mixing ratio, T is the temperature, and c_p is the specific heat at constant pressure and L the latent heat of evaporation. Visibility is diagnosed from these variables and an aerosol content from the UM by first diagnosing the liquid water content, and then using this, or the relative humidity if zero, combined with the aerosol content to estimate visibility. The variational analysis uses forecast first guess values of T_L and q_T together with surface observations of temperature, dewpoint, pressure, visibility and a satellite derived fog mask. The

satellite derived fog is used to provide greater spatial detail than available from the observations alone. For visibilities below 10km T_L and q_T are closely linked to the observed visibility, at higher visibilities this is thought to be inappropriate and they are based on observed temperature and dewpoint. This means that T_L and q_T values are still available when no visibility is observed, thus ensuring maximum use of the available observations. In the forecast, an extrapolation is performed whereby trends of T_L and q_T obtained from the UM are applied to the analysed values. Persistence and the UM forecasts of T_L and q_T are then merged with the extrapolation. The weights for the merging are estimated as a function of the RMSF visibility error against all available observations at T+0 of a forecast from the previous hour. For the objective assessment, analyses and forecasts are compared with observations. Occasions of visibility less than 200m and 1000m are assessed using HR , FAR, and CSI as for cloud. For visibility less than 5000m, RMSF scores are calculated with values greater than 10km limited to 10km to prevent the scores being biased towards high visibilities. Temperature and dewpoint, derived from T_L and q_T, are verified using MEAN and RMS scores. The percentage number of occasions when the forecast visibility is within 30% of the observed and the forecast temperature and dewpoint are within 2 degrees of the observed values are also recorded.

2. Method

To perform the OIS, operational Nimrod software was modified and run on the Nimrod development workstation (NIMDEV) every hour, using observations, satellite data, UM fields and other Nimrod data copied from the operational machine. On NIMDEV two parallel sets of analyses and forecasts were run using specified observation lists over the UK. A 'mainly manual' (MAN) and 'a mainly automatic' (AUTO) list were defined. The observation locations are plotted in figure 2.

These lists were provided by (O) Division and were derived by 'pairing' manual stations with nearby automatics wherever possible to give an approximate 50km average spacing. For verification, three regions of the UK were chosen (figure 3); a complete area with an even coverage of observations referred to later as AREA 1, an area covering northern upland and coastal regions of mainland UK, AREA 2, and remaining part of England, Wales and Northern Ireland, AREA 3. Note that areas 2 and 3 contain a mixture of automatic and manual observations of cloud and visibility and that a few of them do not occur in either of the MAN or AUTO lists. Area 1 also contains a mix of manual and automatic sites but also some ordinary SAMOS and SAWS observations for temperature verification. AREA 1 also includes observations of cloud and visibility not in the MAN or auto lists.

All of the analyses and forecasts using the MAN and AUTO lists were verified against each of the sets of observations comprising areas 1, 2 and 3. It is worth emphasising at this stage that the only difference between the MAN and AUTO runs in terms of data were the different sets of synoptic observations over the UK. Both runs used data from outside the UK (Eire, Channel Islands and the continent) as well as any ship reports. METARS were not used over the UK but were outside (all METAR report identifiers beginning with 'EG' were excluded). Although both the runs used the same UM fields, the T+1 first guess fields would be different, so some 'memory' from the previous analysis was carried forward to the current analysis. This meant that small differences between the MAN and AUTO analyses became apparent where no data were available. However, the greatest differences occured as a consequence of different data availability over the UK. In order to provide a measure of control, the operational Nimrod analyses and forecasts (OP) were verified using the same observation lists as for the experiments. The OP runs used the same UM fields as the experiments but used all observations available, including METARS, over the UK.

As well as the objective verification, selected images for analyses and forecasts were saved for subjective interpretation.

3. Results

The trial was run between 0000 UTC on 1/11/96 and 1200 UTC on 20/11/96. However, there were no results from 0900 on the 10th to 1100 on the 12th because of computer problems. Also, problems with the operational visibility procedures meant that the visibility part of the trial was suspended between 1700 on the 12th and 1200 on the 14th. Further computer problems meant that no images were saved for verification times after 0800 on the 17th.

Unfortunately very little fog occurred during the trial and the only foggy period on the 10th/11th coincided with the period of outage. The period was mainly dominated by cyclonic weather systems giving extensive cloud cover and precipitation at times. The percentage frequency of occurrence in the OIS of critical parameters for cloud and visibility are shown in tables 1 and 2.

Table 1. Percentage frequency of occurrence of cloud parameters in each verification area (not including period 10/0900 - 12/1100 UTC).

AREA	Cloud > 1/2 cover	Cloud > 1/2 cover below 5000 feet	Cloud > 1/2 cover below 1000 feet
1	68	46	7
2	70	42	5
3	68	48	9

Table 2. Percentage frequency of occurrence of visibility parameters in each verification area (not including periods 10/0900 - 12/1100 and 12/1700 - 14/1200 UTC).

AREA	Visibility <= 200m	Visibility <= 1000m
1	0.3	1
2	0.2	1
3	0.4	2

These tables show that during the trial poor visibility occured very rarely and that which occurred was mainly low cloud rather than radiation fog. Also low cloud was relatively rare. The percentages are similar in all areas with very slightly more occasions of poor visibility and low cloud in area 3. The similarity is a consequence of the prevailing weather types during the trial.

3.1 Cloud

During the trial a clear signal began to emerge that for low cloud below 1000 feet, the AUTO run was significantly worse than the MAN and OP runs. However, this was not attributable to differences in quality between the automatic and manual observations but to the fact that none of the automatic cloud observations

were getting into the analyses because of a fault in the HORACE software on the Nimrod computers. Unfortunately this was not found until near the end of the trial. This error had little effect on the MAN runs but a very large effect on the AUTO, increasing the average observation spacing to approximately 100km. To distinguish this from the original specification it is identified as 'AUTO-MAN' in the rest of this section. Thus the available comparisons for cloud are of a reduced manual set of observations MAN, a much reduced set of observations AUTO-MAN and the control OP.

Graphs showing the performance of each run using a variety of objective measures for each critical parameter as a function of lead time in each verification region are shown in figures 4 to 9. The Nimrod results are the graphs headed 'Merged'. Graphs showing the performance of persistence(Per) have been provided for comparison.

The main points from the graphs are:

- For total cover (and also for cover below 5000 feet, not shown), there are only very small differences between the 3 runs.

- There is a strong signal for cloud cover below 1000 feet. There are only small differences between the OP and MAN runs, however, the AUTO-MAN run is significantly worse in all areas, although the differences become smaller with increasing lead times. These are best seen in the graphs for CSI (figures 5,7, and 9).

- For > 3/8 cloud base, MAN is slightly worse than OP and AUTO-MAN performed worst.

- There were only small differences in results between the areas, although in area 1 MAN performed relatively worse than OP.

In order to determine some relationship between network degradation and accuracy of products in Nimrod, the percentage loss of accuracy of the trial analyses and forecasts from the OP runs for significant low cloud cover below 1000 ft in AREA 3 have been tabulated in table 3. This was achieved by subtracting the trial CSI score from the OP score for the corresponding lead time and then dividing this by the OP score. The MAN network has an approximate observation spacing of 50km in this area and the AUTO-MAN 100km compared with an OP value of about 45km. The UM scores have also been provided to give some idea of the possible accuracy if no surface observations were used in the Nimrod analyses except to provide T+0 fields for the UM at 00, 06, 12 and 18 hours. Since the UM performance only falls off slowly with lead time, an average of scores for T+4 to T+9 forecasts have been used, since it is these lead times that are used in the cloud merging in Nimrod.

Table 3. Percentage reduction in accuracy of > 4/8 cloud cover below 1000ft in the OIS verified against observations over England, Wales and Northern Ireland (Area 3) using CSI. The two experimental runs had manual observation spacings of 50km (MAN) and at least 100km (AUTO-MAN) and were compared against the operational Nimrod run using all manual observations. The UM results are based on average scores of operational UM T+4 to T+9 forecasts from model data times.

	MAN (50km)	AUTO-MAN (>100km)	UM
Nimrod (merged) T+0	0	39	79
Persistence T+1	4	32	70
T+2	2	28	63
T+3	3	27	57
T+4	3	23	48
T+5	0	22	41
T+6	4	24	36
Nimrod (merged) T+1	0	36	70
T+2	0	30	60
T+3	0	28	50
T+4	0	22	41
T+5	0	17	30
T+6	0	23	27

The largest degradation is in the Nimrod (merged) analyses. For persistence there is a small percentage reduction between MAN and OP, but a relatively large one between AUTO-MAN and OP decreasing from 32% at T+1 to 24% by T+6. The reason for the reduction is that OP persistence forecasts are much more accurate at T+1 than T+6, so the impact of degradation decreases with lead time. The UM reduction also gets smaller because the UM performance remains constant while the OP Nimrod gets worse. In the Nimrod (merged) forecasts, the degradation between OP and MAN is negligible since in producing it's forecasts Nimrod uses weights based on the recent accuracy of persistence, extrapolation and UM forecasts. This means that less weight would have been given to MAN persistence and extrapolation than in corresponding OP runs making up for any potential loss of accuracy. However, in the AUTO-MAN, persistence and extrapolation would have performed much worse than in MAN thus more weight would be given to the UM than in the OP run. This is evident in the greater reduction in accuracy at T+1 to T+3 than persistence. After T+3, more weight is given to the UM in OP and so AUTO-MAN Nimrod starts to perform better than AUTO-MAN persistence, and the UM reduction decreases as the OP Nimrod becomes more like the UM forecast.

The average of the T+1 to T+6 accuracy reductions in Nimrod is 26% which means that by degrading the observing network from 50 to 100km spacing we can expect an average degradation in performance of 26%.

These results obviously must be treated with caution because of the limited amount of data, however, they do suggest that for low cloud forecasting, small (less than 5%) increases in observation spacings would have negligible impact on average but that an increasing loss of accuracy can be expected with cuts above 5%.

An example of the effect of using a degraded set of observations is illustrated in the following example from the trial shown in figure 10. At 0000 UTC on 6/11/96, the low cloud cover in MAN is very similar to OP with all the main cloudy areas analysed but small differences around observations not used in MAN. However, the AUTO-MAN analysis has many cloudy regions missing, notably around Southampton, the Midlands, Cornwall, Wales, northeast England and northern Scotland. All these regions had no manual observations in the AUTO-MAN run. For total cover (figure 11), there are only small differences between the runs which is a consequence of combining observations with satellite and model information to provide the analysis. In predominantly cloudy conditions surface observations provide detail of the lowest layers and the model and satellite of the higher layers.

3.2 Visibility

Graphs showing the performance of each parameter are shown in figures 12-17 .The results do not include verifications between 12/1800 and 14/1200 because of the missing OP runs during this period. Comparing the observation availability between the MAN and AUTO (figure 2), we see that the MAN network has gaps over the Scottish Highlands, Lancashire and Wales. The AUTO network also has a gap over Lancashire as well as the northwest Midlands (MAN has only one observation here) and W. Scotland. Because of the relatively few occasions of visibility below 200m, this report will concentrate on visibility below 1000m threshold and the other parameters.

Looking first at temperature and dewpoint (figures 13,15, and 17): in all areas we see an almost linear decrease of accuracy with lead time, with Nimrod (merged) performing better than persistence. The operational temperature and dewpoint forecasts are very good with more than 85% within 2 degrees of observed values at T+6. Both the MAN and AUTO performed worse than OP. However, it was the AUTO run which was worse in area 1, the MAN in area 2, and MAN in area 3. For temperature, MAN was almost the same as AUTO, (figure 17). This variability according to area occurred also with the other parameters.

For visibility below 5000m (figures 13,15 and 17) the operational results show a steady increase in RMSF error with lead time. The low percentage of occasions when Nimrod is within 30% of observed values reflects the natural variability of visibility. In area 2 the results for RMSF are worse by almost a factor of 2 than those in areas 1 and 3. This result probably reflects the greater variability and hence the greater difficulty of forecasting visibility in the northern upland regions. Both the MAN and AUTO runs were significantly worse than OP with AUTO worse in area 1 and MAN worse in areas 2 and 3.

For visibilities below 1000m the variability between areas is similar to temperature, which is not suprising as the temperature and dewpoint response is going to be greatest when there is fog. However, a clear signal emerges on the impact of observations on forecasts in this visibility range. Looking at the CSI scores (figures 12,14, and 16), there is very little difference between the runs after T+3, however, before then the trial runs are worse than the operational although all the scores are poor. In area 1 a high FAR is the main contributor to the poor performance of AUTO (figure 12). In area 2 (figure 14) MAN Nimrod is slightly worse than AUTO in the early stages of the forecast due to a slightly lower HR and higher FAR.

The general conclusion that can be made from these results is that degrading the observing network has a detrimental effect on Nimrod forecasts of visibility and

temperature, particularly for poor visibility at ranges up to T+3 hours. Also the different results in each area between MAN and AUTO suggests that the distribution of the observing stations is important, and perhaps more important than small differences in quality between automatic and manual observations. In area 2 the T+1 AUTO forecasts were probably more representative of conditions on the higher ground than MAN.

The importance of a good distribution of observations is illustrated in the following examples.

Figure 18 shows visibility analyses from the MAN, AUTO and OP runs for the 8th at 0600 UTC. Over England there are large differences in the analyses. Charts of plotted observations are shown in figures 19 and 20 which also indicate the observations used in the MAN and AUTO runs respectively. The main foggy region was around Hurn airport, labelled 'H' in figure 19, which had a visibility of 800m in fog with a temperature and dewpoint of -2C. This area extended north and east with Boscombe Down 'B' reporting thinning fog at 1000m and Southampton observing adjacent fog. The operational analysis picked out this area very well whereas both the MAN and AUTO runs did not. Hurn on this occasion was obviously a vital observation and was not used in either the MAN or AUTO runs. The MAN run used Southampton but not Boscombe Down and vice-versa for AUTO. Neither observation was sufficient to analyse an area of fog. Hurn is a well known frost hollow in southern England, being on sandy soil, and is prone to fog development. Hurn was also one of the verification stations used in area 1 but not in area 3.

There were other differences elsewhere, particularly in the Midlands. Both the MAN and AUTO analyses have regions where the visibility is too poor. Inspection of the plotted charts shows that these occur where ther are no observations of good visibility available in the respective runs. The analyses in these regions consequently look very similar to the corresponding first guess forecasts shown in figure 21.

Figures 22 and 23 show another example for the 15th at 0900 UTC. The largest difference between the runs occured over SW England. The MAN analysis has a large area of fog with visibility below 1000m, whereas the AUTO run was clear, and OP had patchy fog. Figure 23 shows that the first guess for all runs had good visibility (more than 5000m), so the observations had a large influence. The plotted chart showing the observations available to the MAN and AUTO runs over SW England is shown in figure 24. In the AUTO analysis the lowest visibility observed was 4400m at Dunkeswell, however, the observations at Exton and St Mawgan in the MAN analysis both had visibilities of 700m with no other observations in between. Only the OP run had the observations at Chivenor, Hartland Point, and Bastreet which gave visibilities around 5000m. Consequently the OP run depicted the patchy nature of the fog best of all. Other differences occured over Wales and northern England where the automatic observations were better placed to pick out the generally better visibilities over high ground.

4. Conclusions

No conclusions can be made from this study on the relative merits of automatic versus manual observations for cloud analysis and forecasting in Nimrod. However, it has been demonstrated that removing large numbers of surface observations has a detrimental impact on the analysis and prediction of low cloud. It must also be emphasised that this impact was seen mainly on low ground as most upland stations (being automatic) were not available to either analysis or verification.

For visibility the results were more complex but the evidence from this study suggests that an even and more dense network than 50km is required for accurate analysis and prediction. This is based on the evidence from the selected cases showing the high variability of visibility over small areas, and the objective verification which showed the operational and more dense network performing better. The results also indicate that an even coverage of observations is

probably more important than differences in accuracy between manual and automatic stations. This is evident from the results in area 2 which showed the AUTO performing better than MAN.

Because of the link between visibility, temperature and humidity in Nimrod it has been shown that degradation in visibility forecasts is matched by a corresponding degradation in temperature and dewpoint forecasts, which in Nimrod are usually very accurate when the diurnal variation is small.

More generally, much experience was gained in conducting observation impact studies within Nimrod. It has become clear that improvements are required in monitoring observations in the system. Also with the visibility component still under operational trial, a lot of effort was required to maintain compatability of the trial runs with the operational owing to code changes. This need could have been prevented by re-running the operational version as well as the experiments. However, there was not enough time in the hourly re-run cycle to do this. Similar studies should only be conducted when the software is stable.

This study did not use the modified cloud analyses to re-initialise the UM, thus the same model information was used for the trial runs. It is expected that feeding back the modified analyses and re-running the UM in parallel would probably have degraded the results for MAN and AUTO more than presented here.

5. Recommendations

- Another impact study should be performed to determine likely effects of automation on cloud forecasts. Comparing runs using only automatic or only manual observations will not be sufficient as it has been shown in this study that degrading the network in terms of observation spacing reduces forecast quality. A possible way may be to add appropriate random noise to all manual cloud observations. The amplitude of the noise should be consistent with expected errors from automatic observations and may need to be varied according to cloud base.

- A similar approach for visibility may also be appropriate. However, it is clear that the optimum affordable observation spacing for visibility forecasting in Nimrod needs to be found. Putting noise into selected manual observations and using all the automatic observations would give a measure of the impact of automation with the present network. However, field experiments in selected areas using a more dense network of visibility and temperature observations would have to be undertaken to determine additional benefits of a denser automatic network.

- There were few occasions of formation and clearance of dense fog in this study and observation impacts in such conditions remain to be assessed. To address this, an OIS could be set up 'ready to go' when foggy conditions are anticipated in the medium range forecasts, thus building up a collection of statistics over a period of time. This would, however, be dependent on loading of the Nimrod development computer. An alternative would be to perform case studies, however, assembling the relevant data would be time-consuming and probably impractical. In dense fog conditions it would be worth considering impacts of observation spacings to find a relationship between forecast accuracy and observation spacing for visibility.

- In future experiments the OP version should be re-run to guarantee 'clean' comparisons. This implies that it may not be possible to do cloud and visibility studies concurrently.

- Cross-comparison of OIS results with Nimrod analyses should be considered, particularly for fog.

6. References

Golding, B.W. (1996), "Nimrod: a system for generating automated very-short-range forecasts". Submitted to Met. Apps.

Hand, W.H. (1995), "Report on test of merging extrapolation and Mesoscale model cloud forecasts for Nimrod". Forecasting Research Division Technical Report No. 163.

Wright, B. and Thomas, N. (1996). "An objective Visibility Analysis and Very-Short-Range Forecasting System". Forecasting Research Division Technical Report No. 200.

7. Acknowledgements

The author would like to record the work done by Derek Steadman in Forecasting Products in writing the control code and monitoring the OIS experiments, and also for providing the objective performance graphs.

Also Eddie Spackman in (O) Division for providing the figures showing the locations of the observing stations in the experiments.

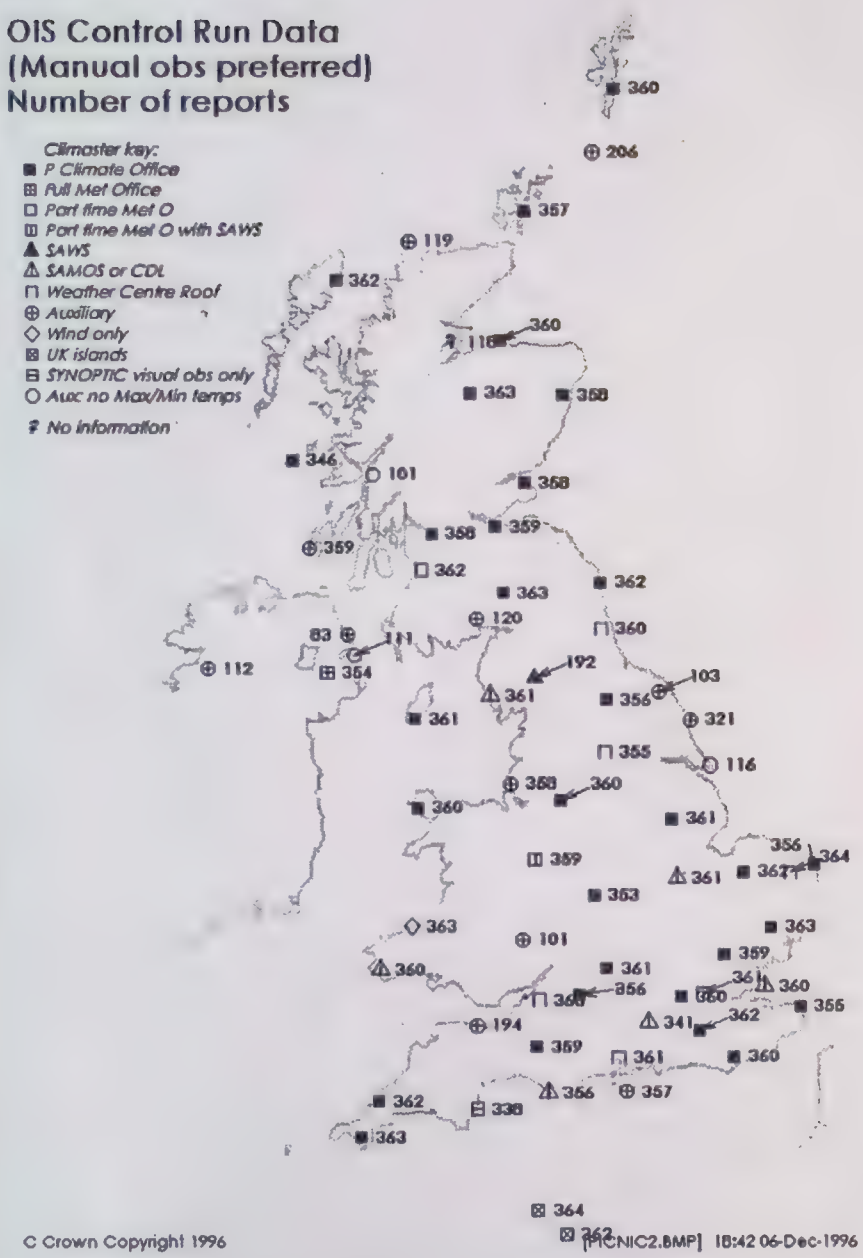

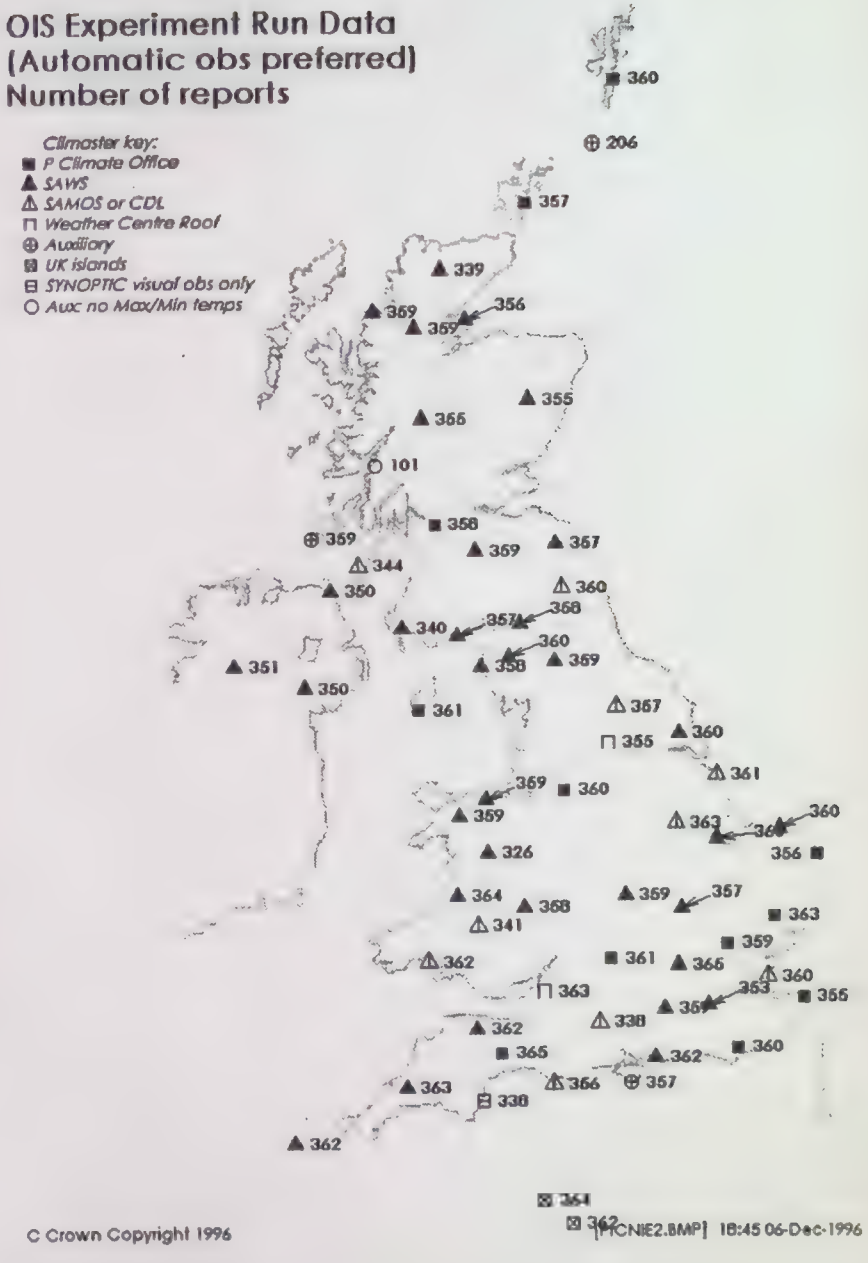

Figure 2. Plots of observing stations used in the MAN (left) and AUTO (right) runs. Triangles indicate completely automatic stations. The numbers indicate the number of occasions when the observation was available to the OIS runs.

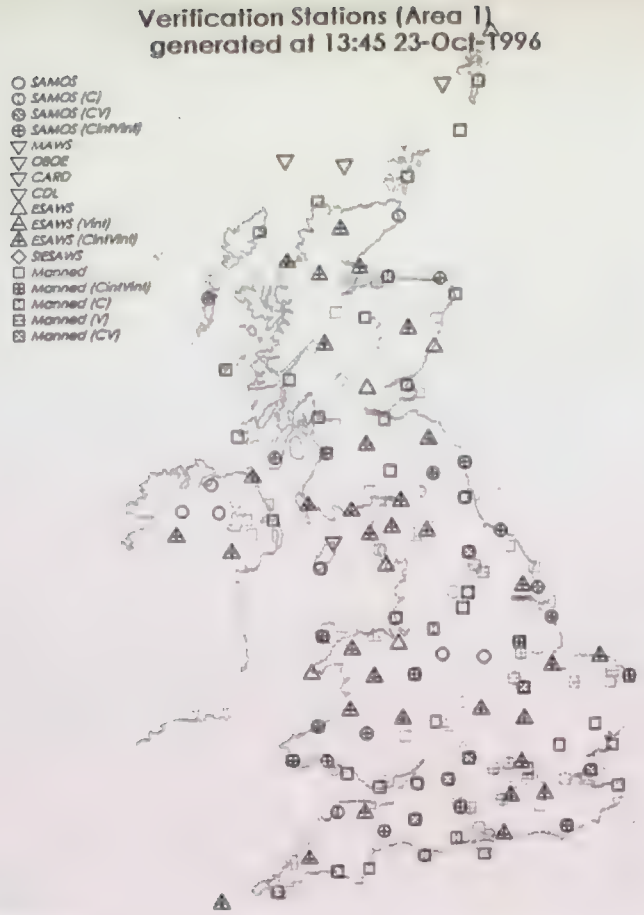

Figure 3. Plots showing the location of stations used for verification in each of areas 1,2 and 3.

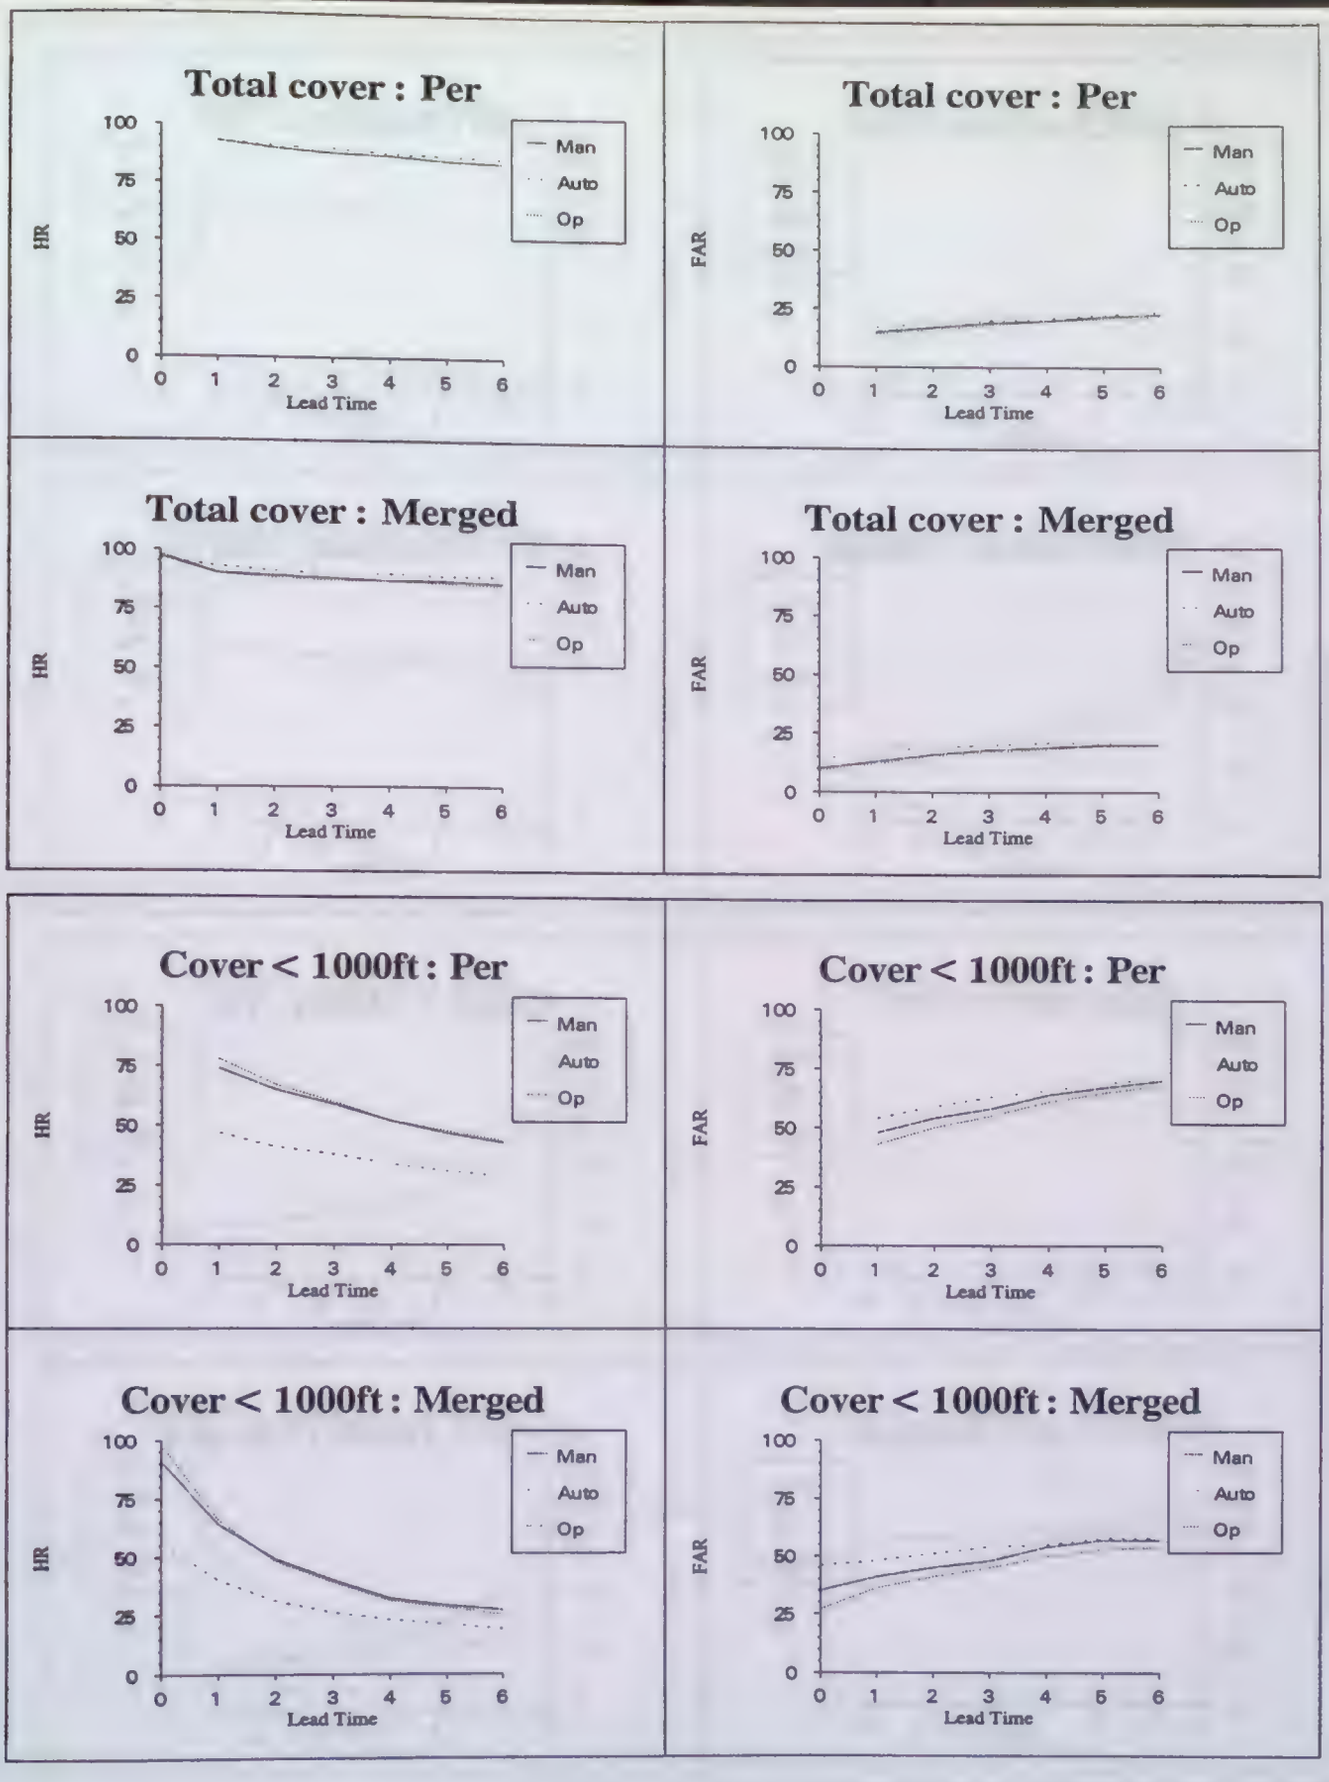

Figure 4. Graphs showing the performance of the operational and test runs for > 4/8 total cloud cover and > 4/8 cloud cover below 1000 feet in area 1. Hit rate (HR) and false alarm rate (FAR) percentages as a function of lead time are shown for persistence and the Nimrod (merged) analyses and forecasts.

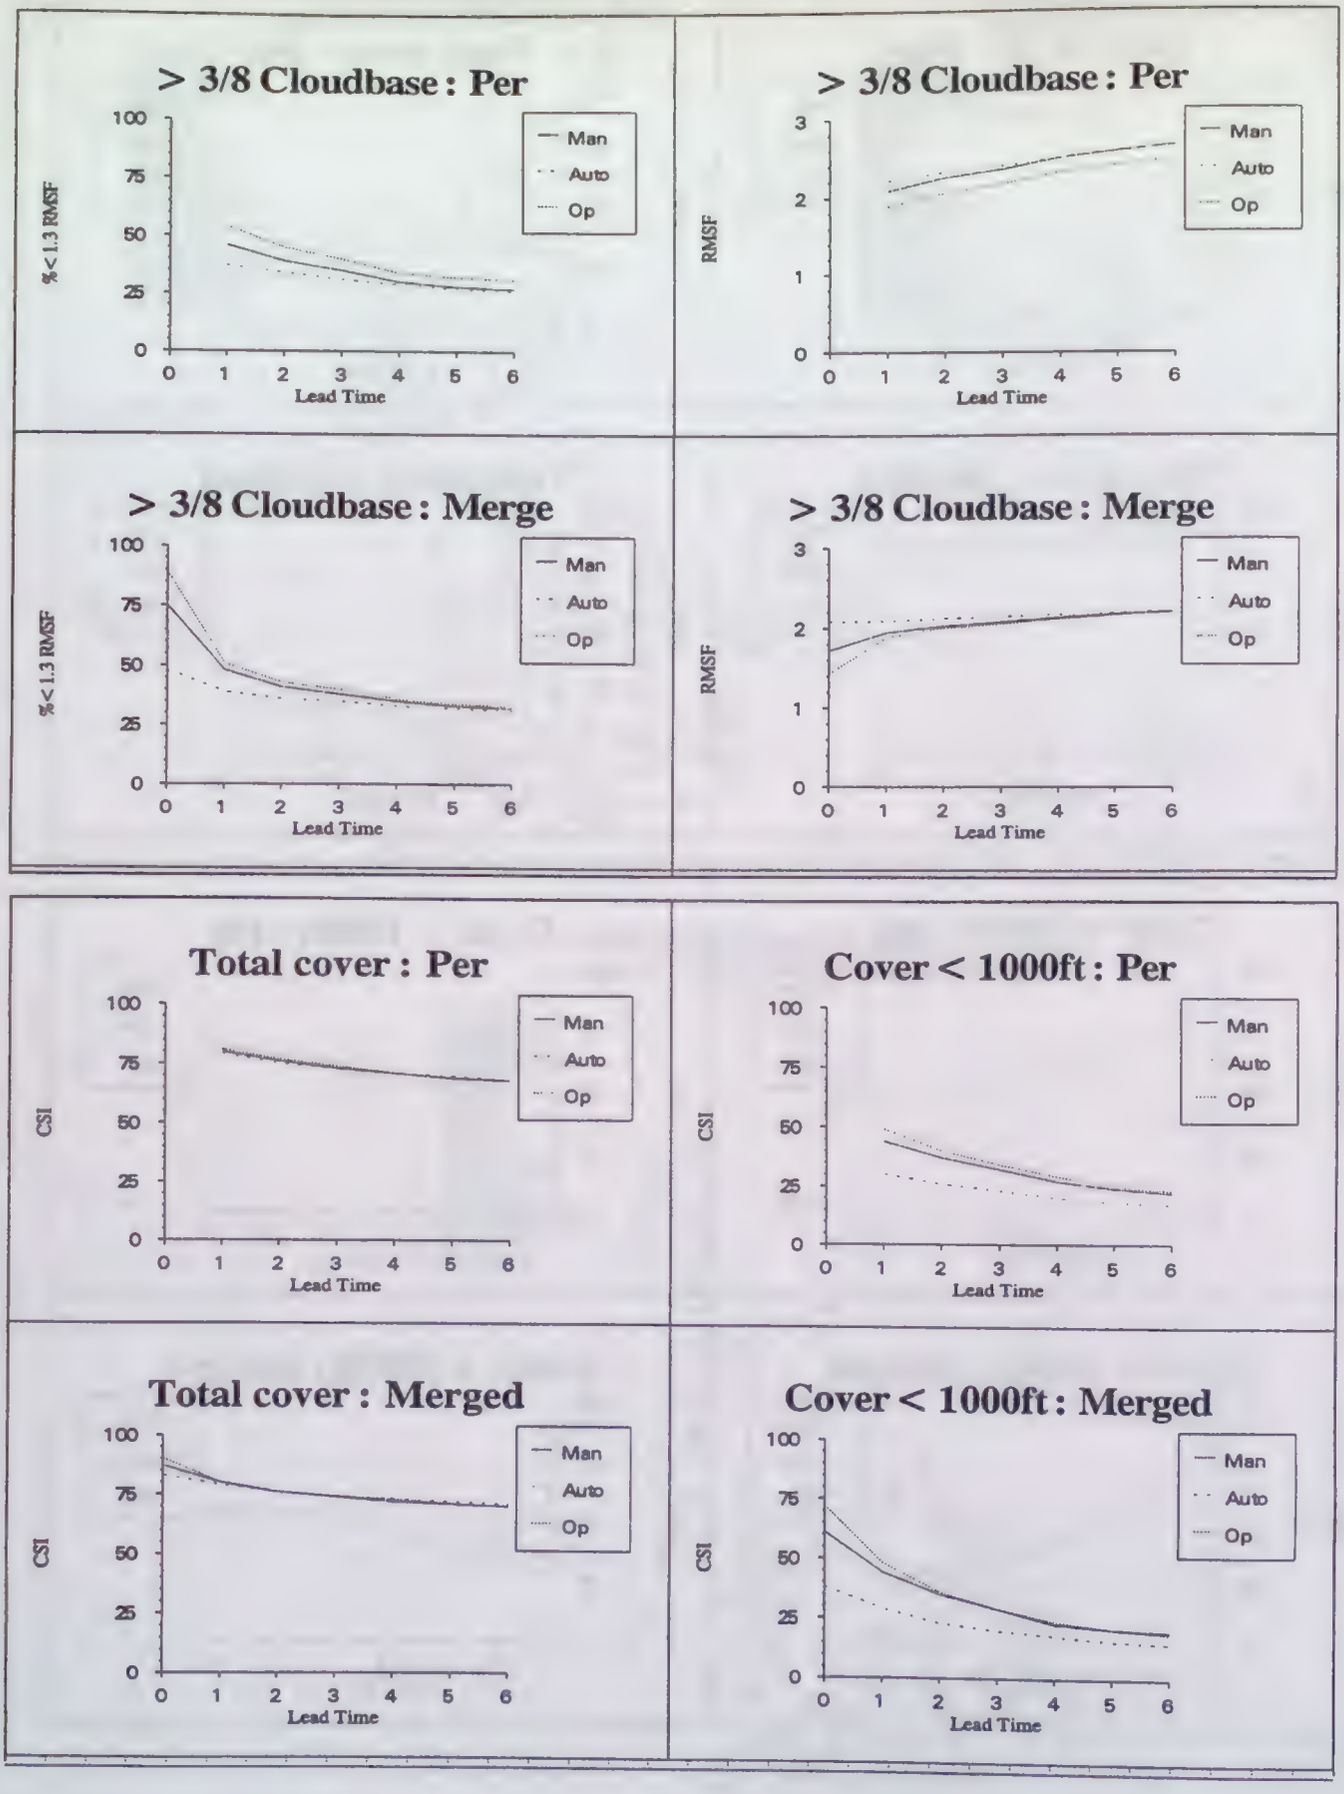

Figure 5. Graphs showing the performance of the operational and test runs for persistence and the Nimrod (merged) analyses and forecasts in area 1. The top graph shows >= 3/8 cloud base. RMSF and the percentage of occasions when the visibility is within 30% of the observed value (RMSF < 1.3) are plotted as a function of lead time. The bottom graph shows CSI for total cover and cover below 1000 feet.

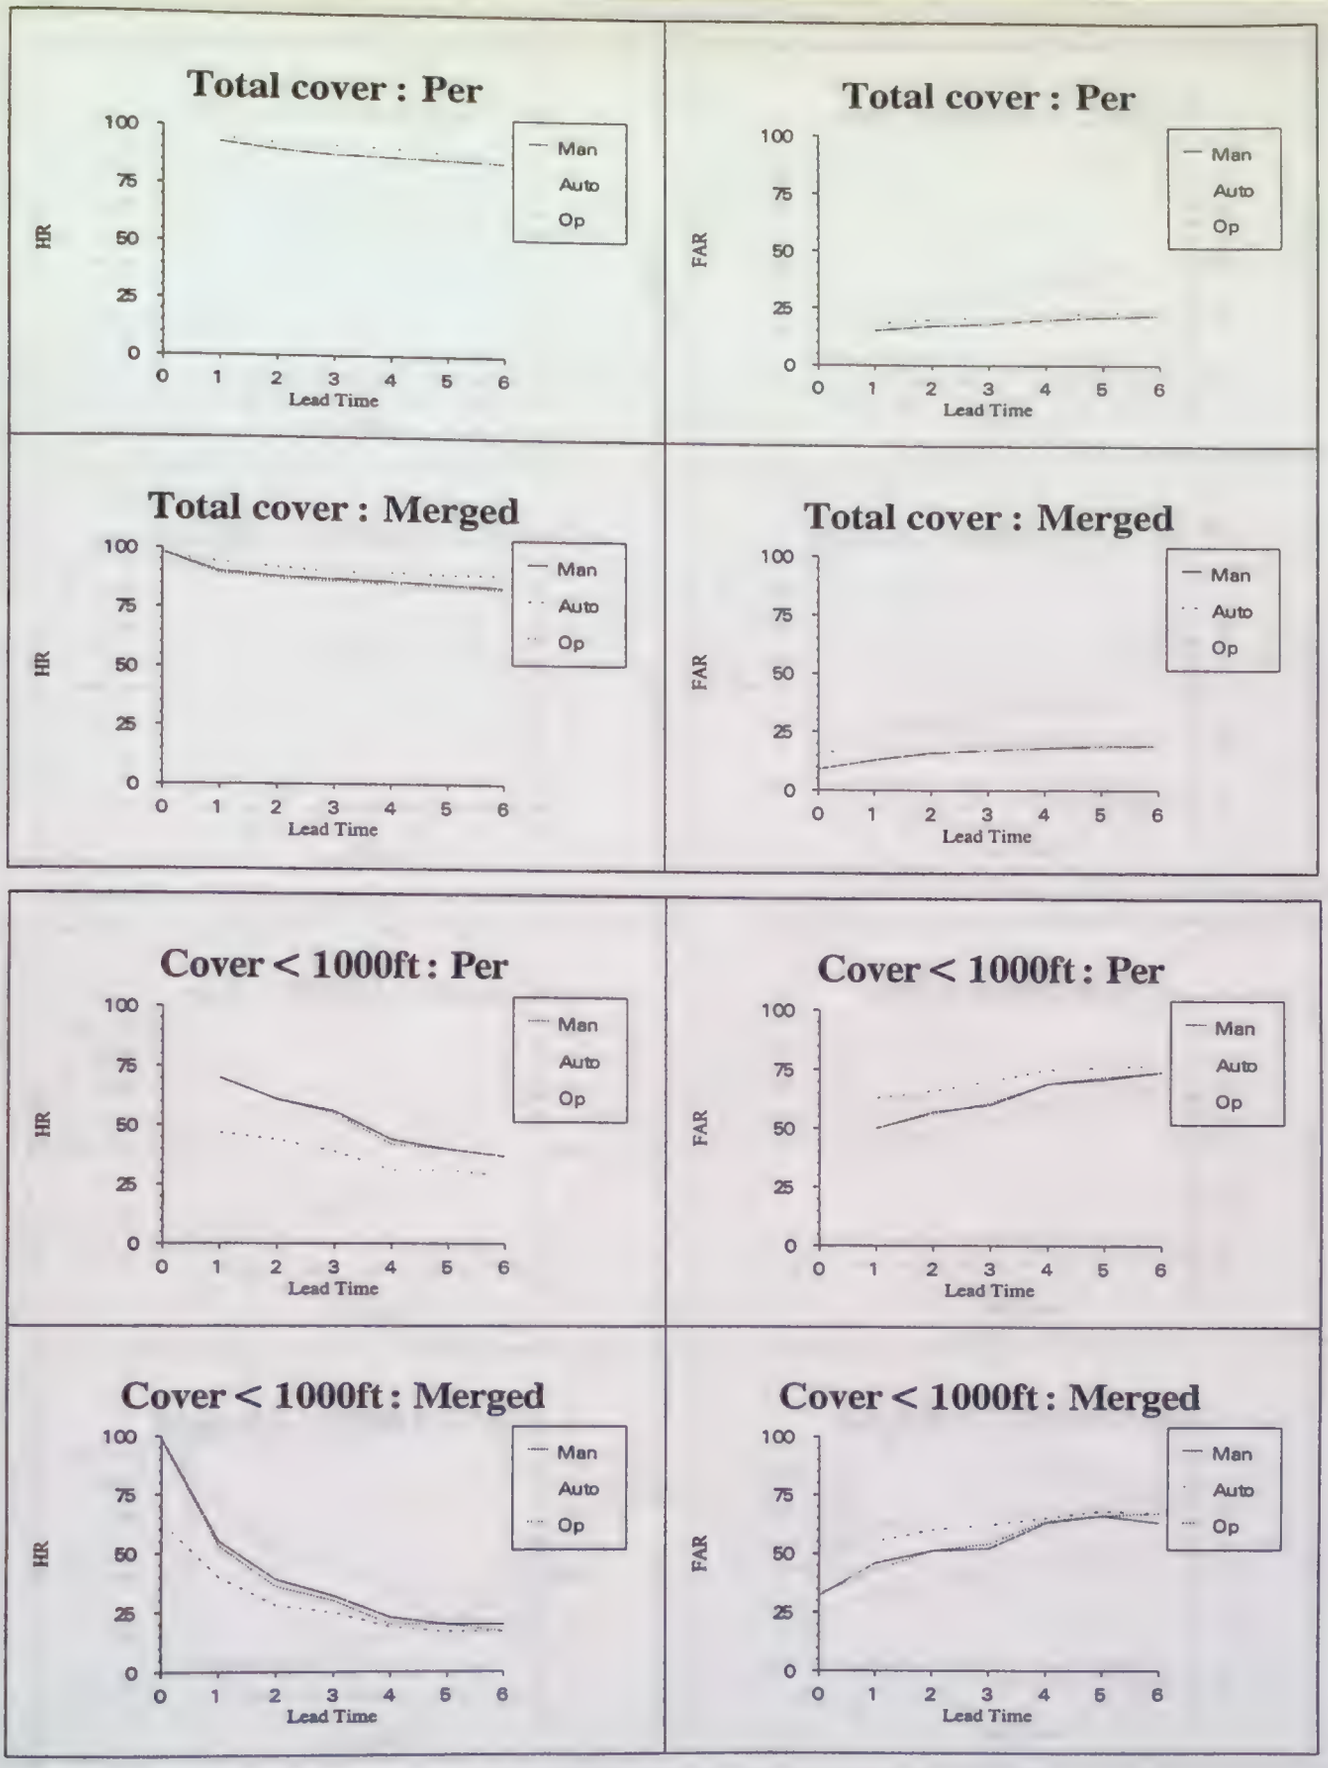

Figure 6. Graphs showing the performance of the operational and test runs for > 4/8 total cloud cover and > 4/8 cloud cover below 1000 feet in area 2. Hit rate (HR) and false alarm rate (FAR) percentages as a function of lead time are shown for persistence and the Nimrod (merged) analyses and forecasts.

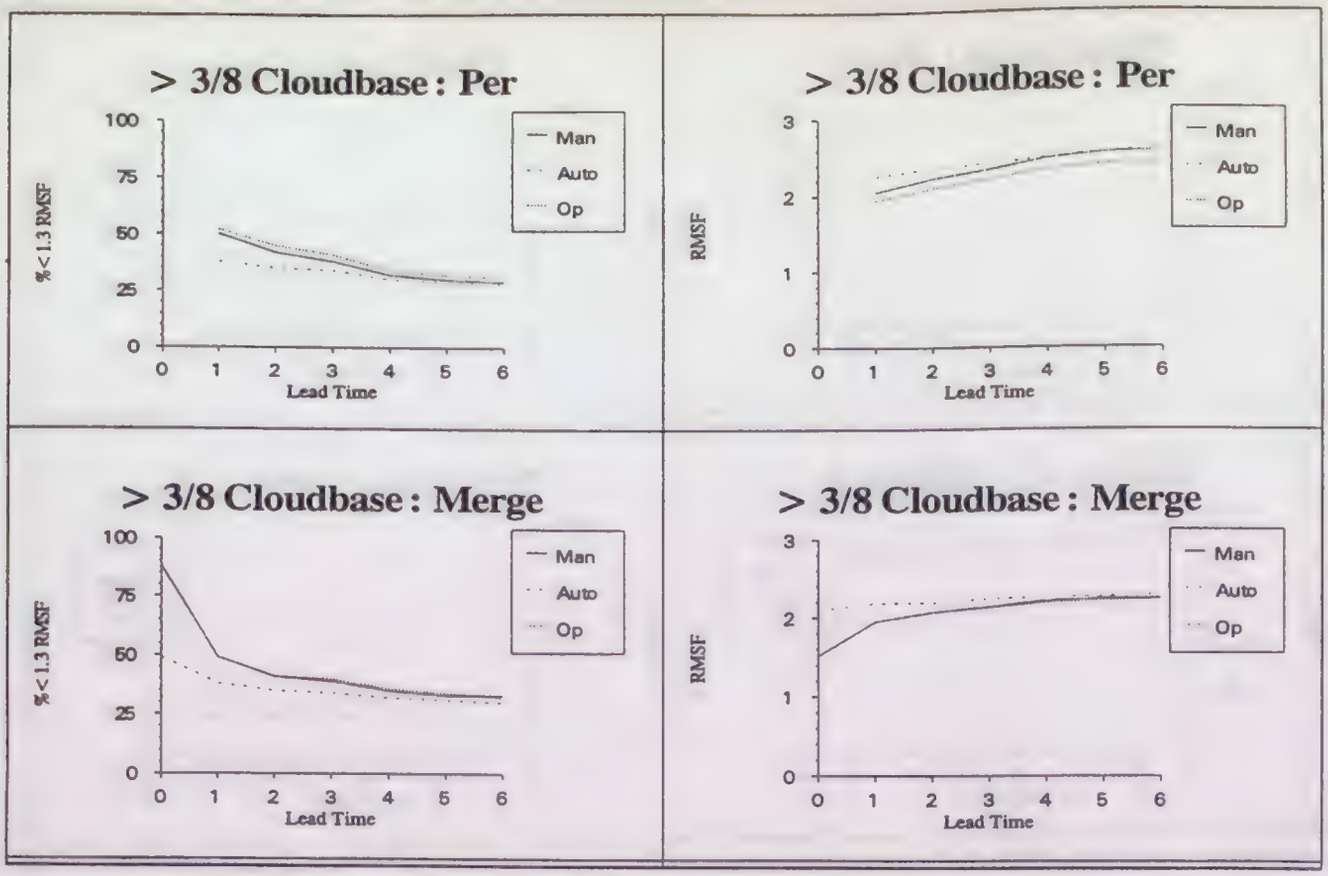

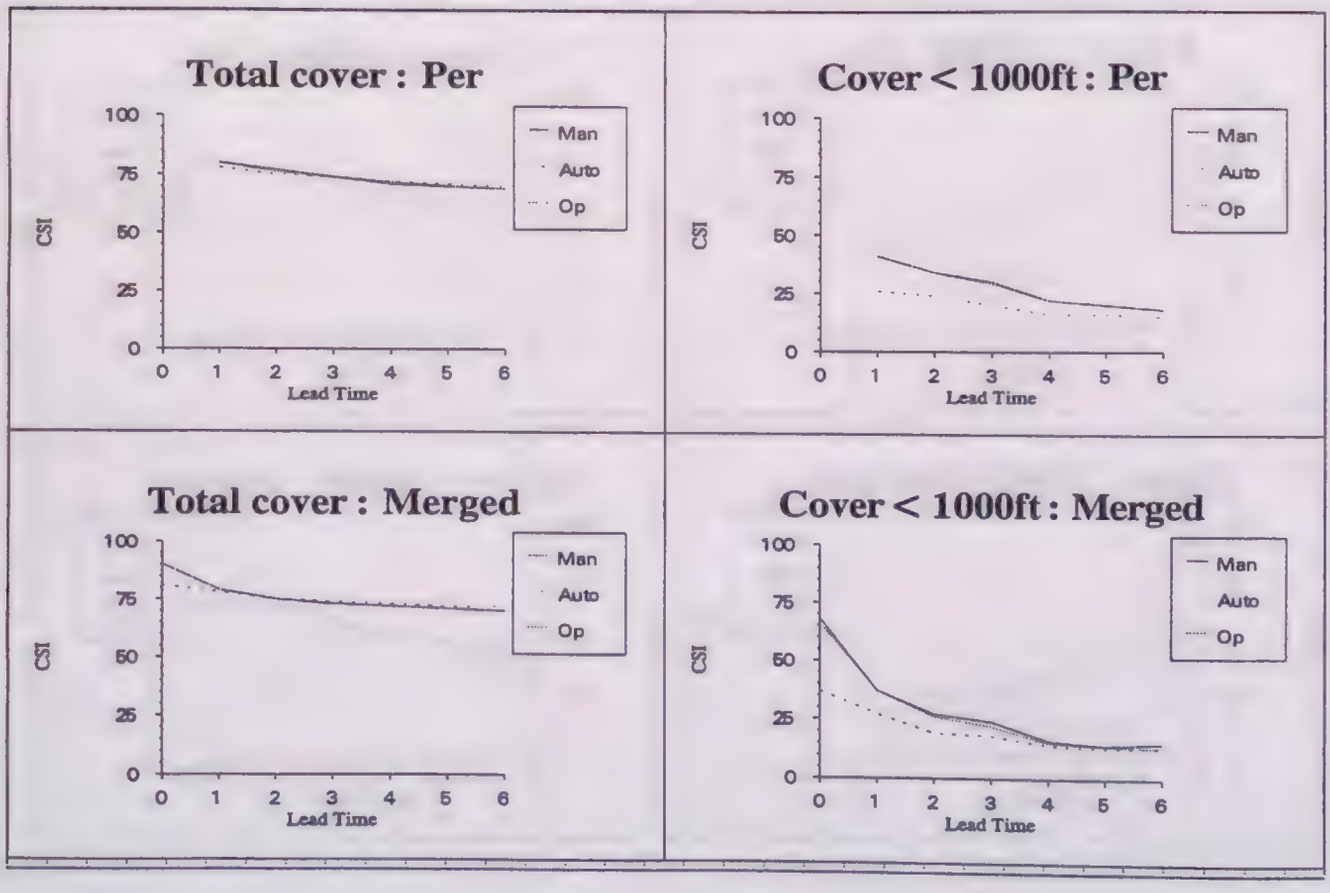

Figure 7. Graphs showing the performance of the operational and test runs for persistence and the Nimrod (merged) analyses and forecasts in area 2. The top graph shows >= 3/8 cloud base. RMSF and the percentage of occasions when the visibility is within 30% of the observed value (RMSF < 1.3) are plotted as a function of lead time. The bottom graph shows CSI for total cover and cover below 1000 feet.

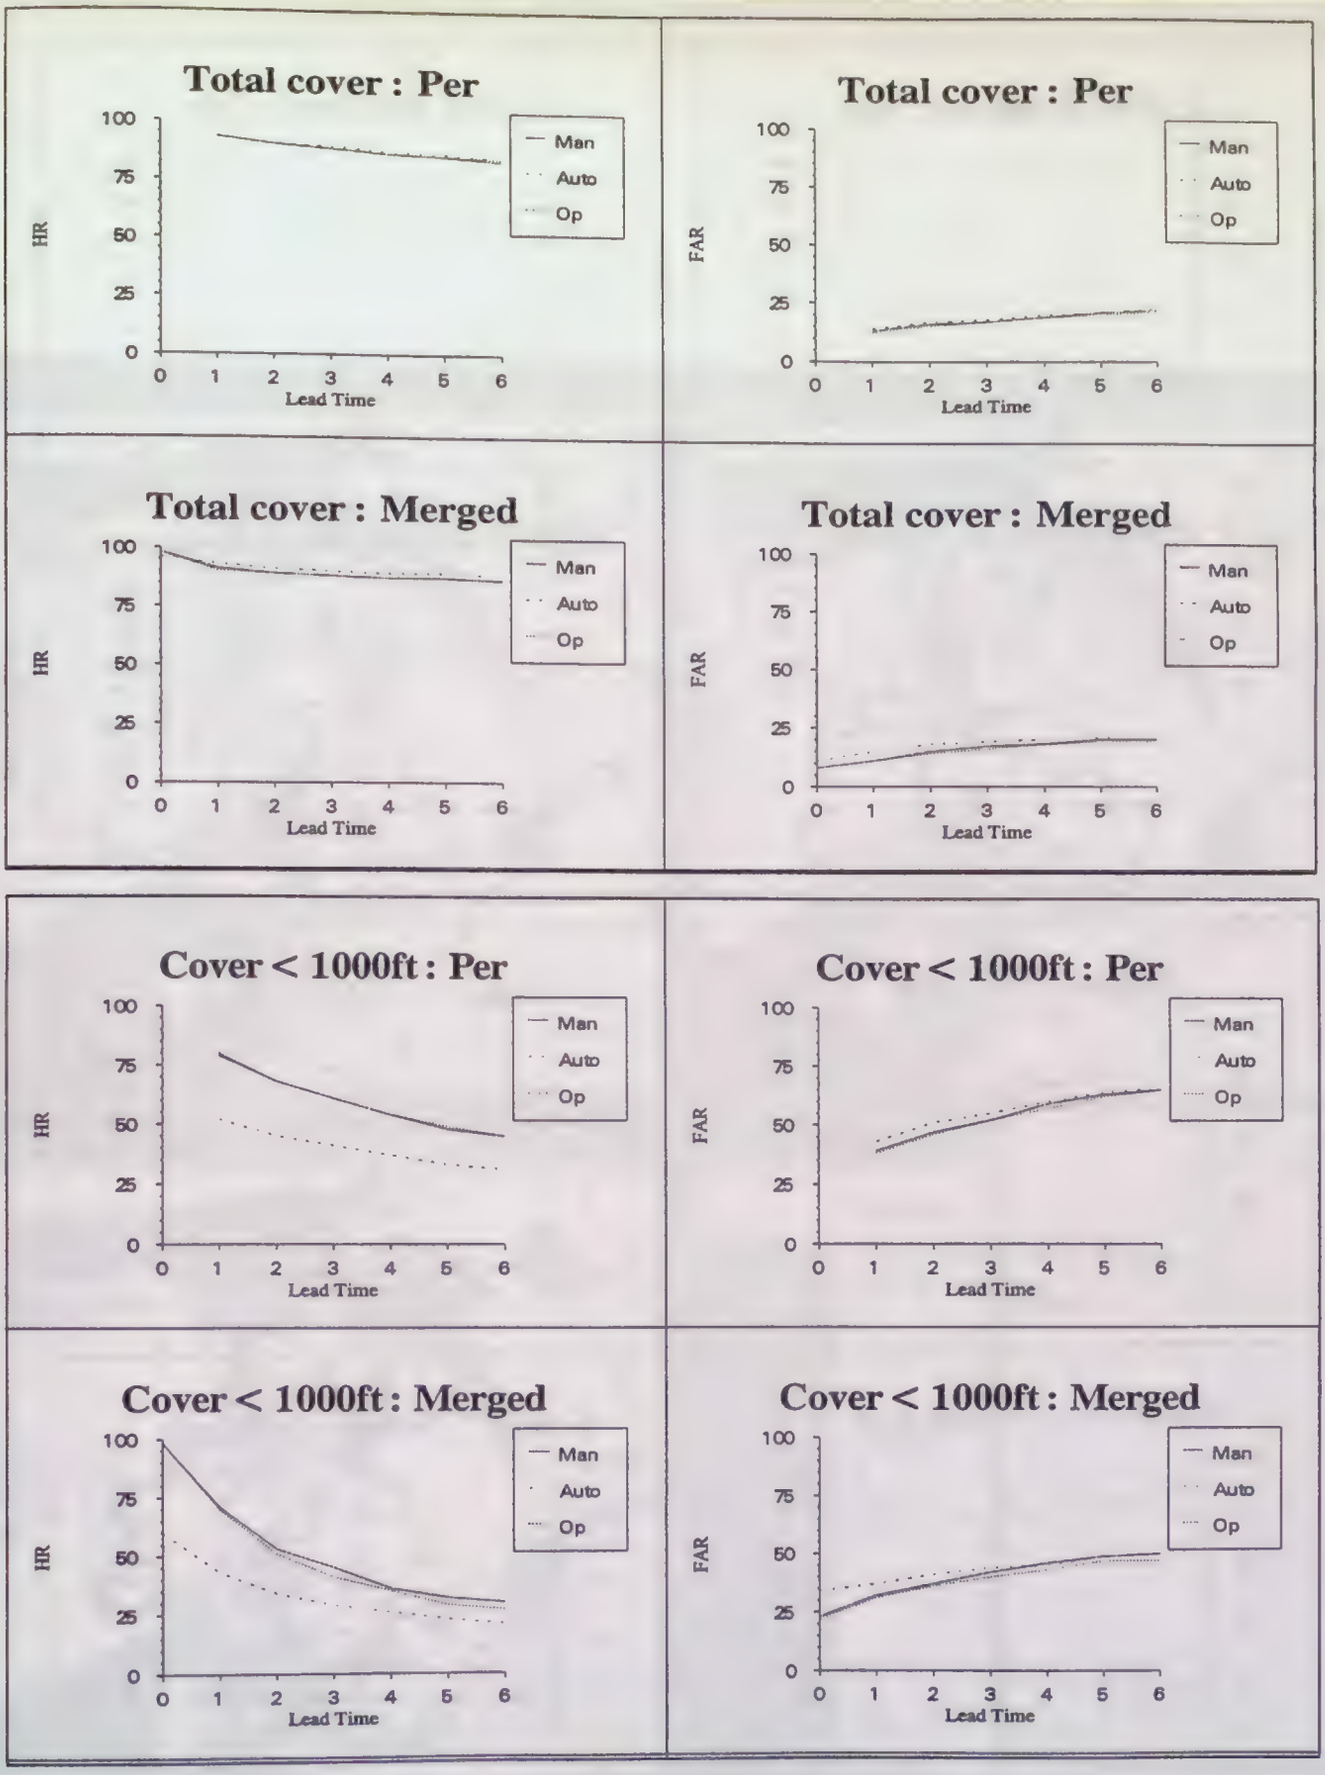

Figure 8. Graphs showing the performance of the operational and test runs for > 4/8 total cloud cover and > 4/8 cloud cover below 1000 feet in area 3. Hit rate (HR) and false alarm rate (FAR) percentages as a function of lead time are shown for persistence and the Nimrod (merged) analyses and forecasts.

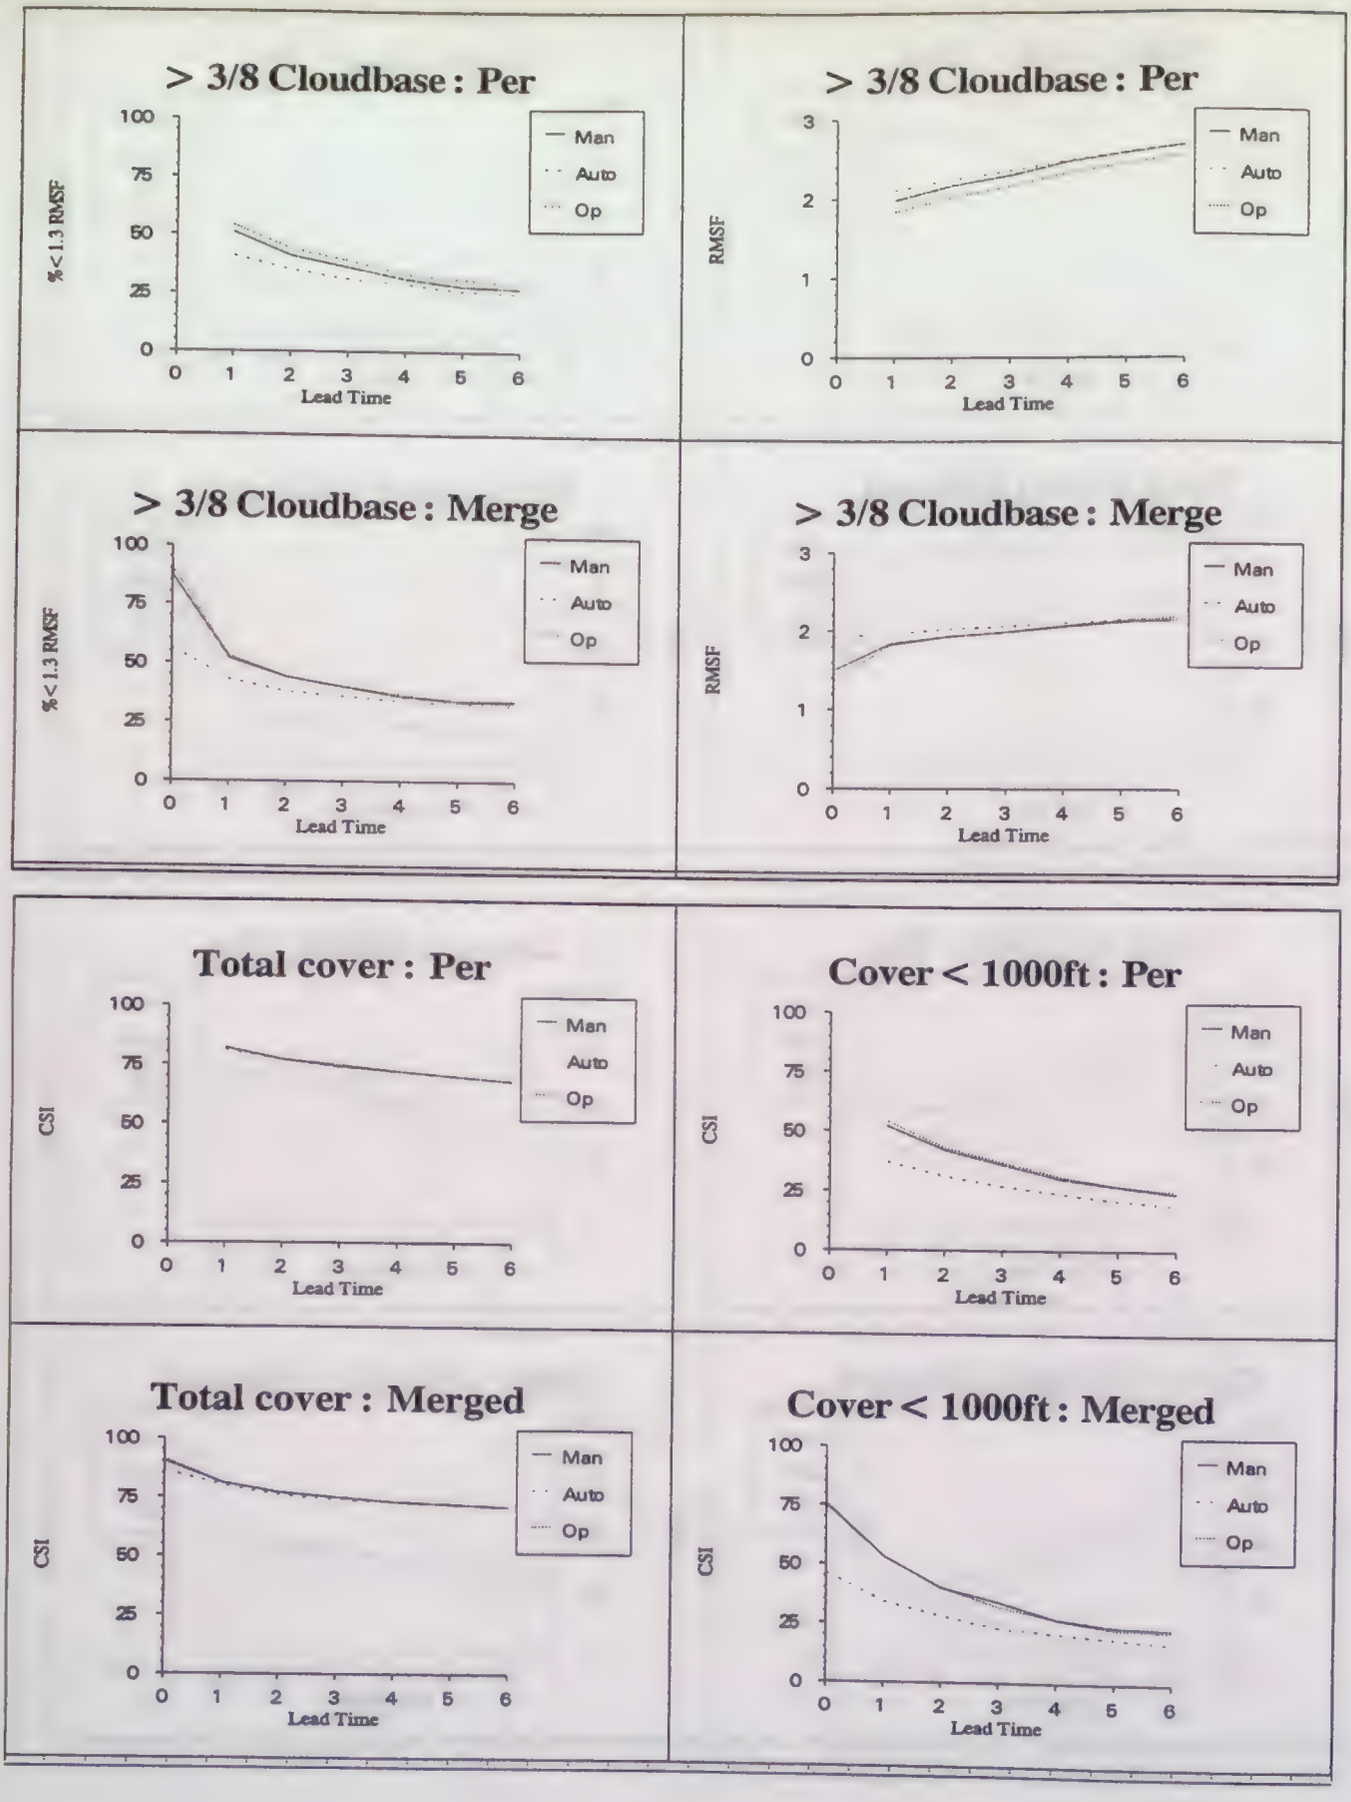

Figure 9. Graphs showing the performance of the operational and test runs for persistence and the Nimrod (merged) analyses and forecasts in area 3. The top graph shows >= 3/8 cloud base. RMSF and the percentage of occasions when the base is within 30% of the observed value (RMSF < 1.3) are plotted as a function of lead time. The bottom graph shows CSI for total cover and cover below 1000 feet.

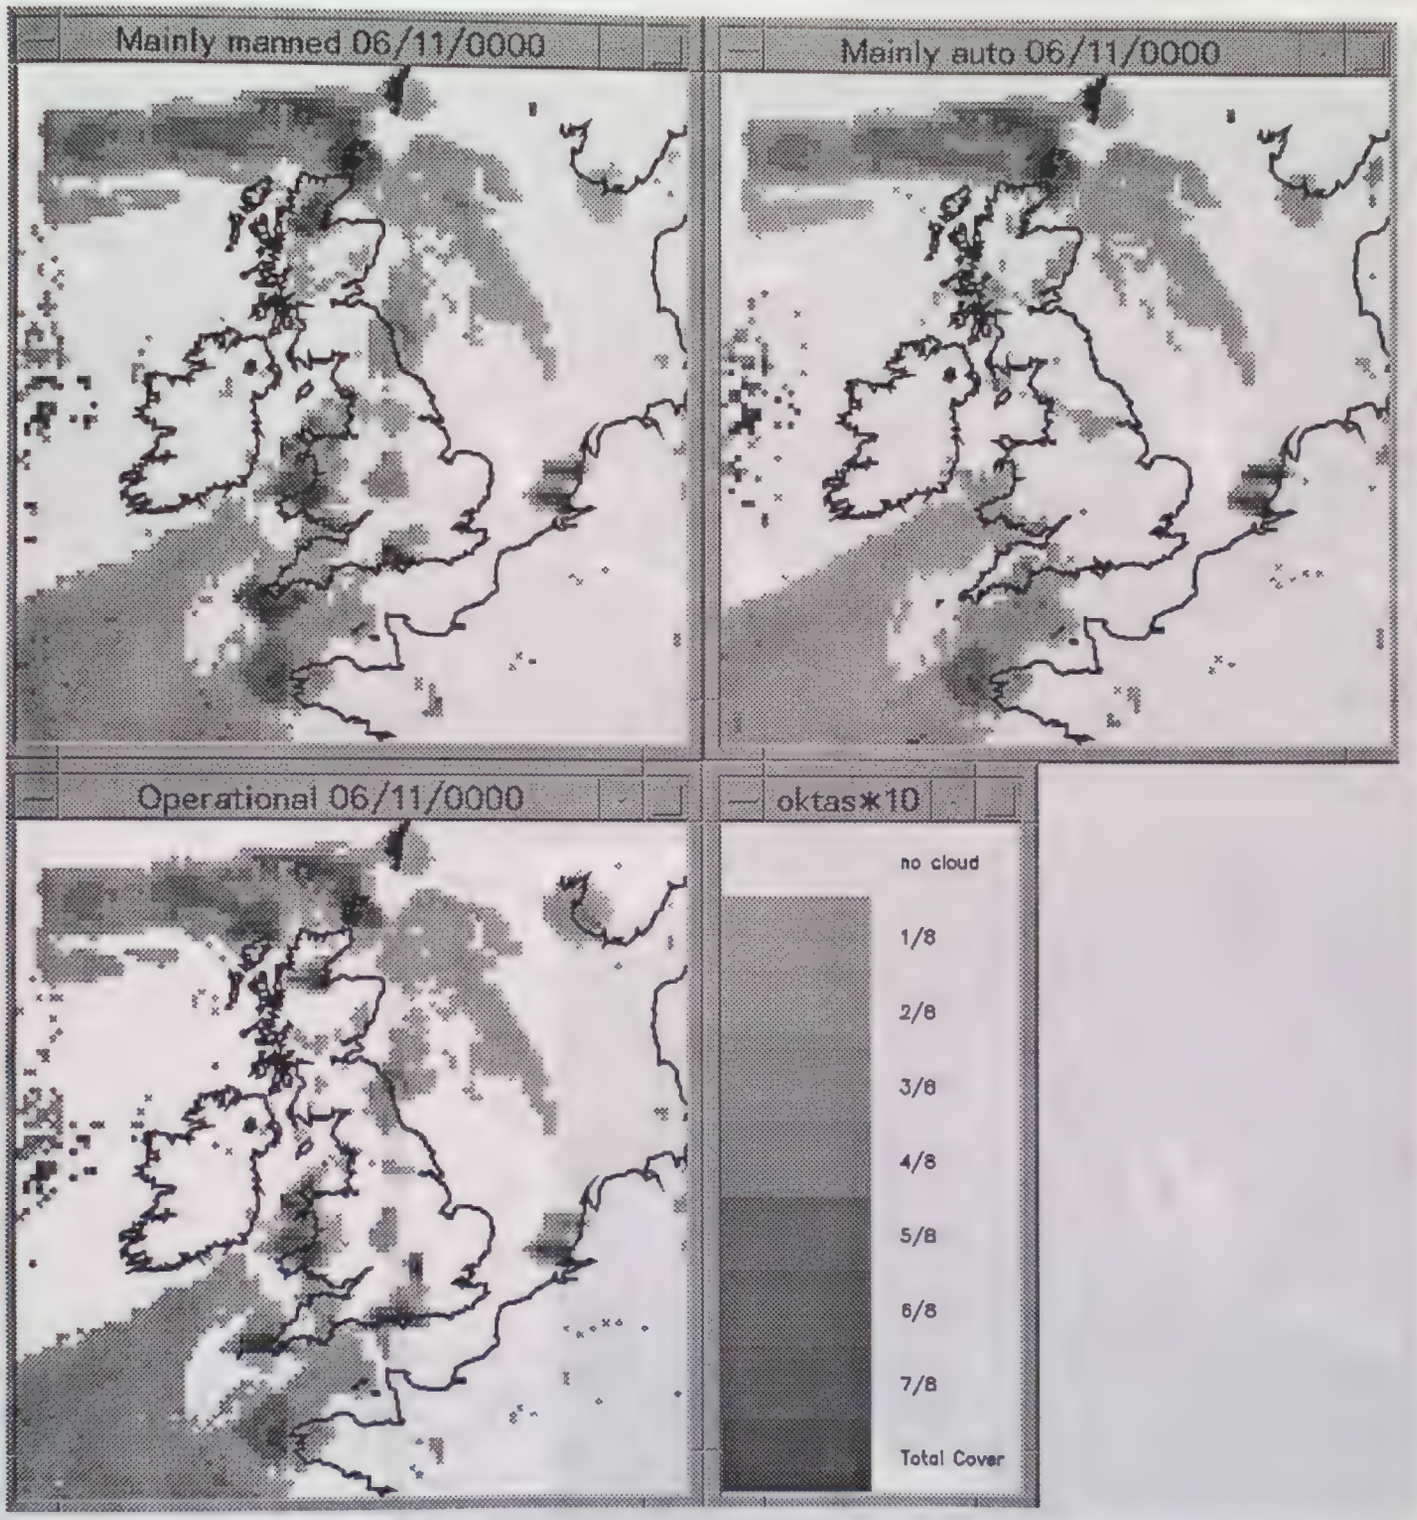

Figure 10. Charts showing amount of cloud cover below 1000ft in each 15km pixel from the MAN, AUTO-MAN and OP analyses for 0000 UTC 6/11/96. Cloud amounts are shaded according to coverage in oktas.

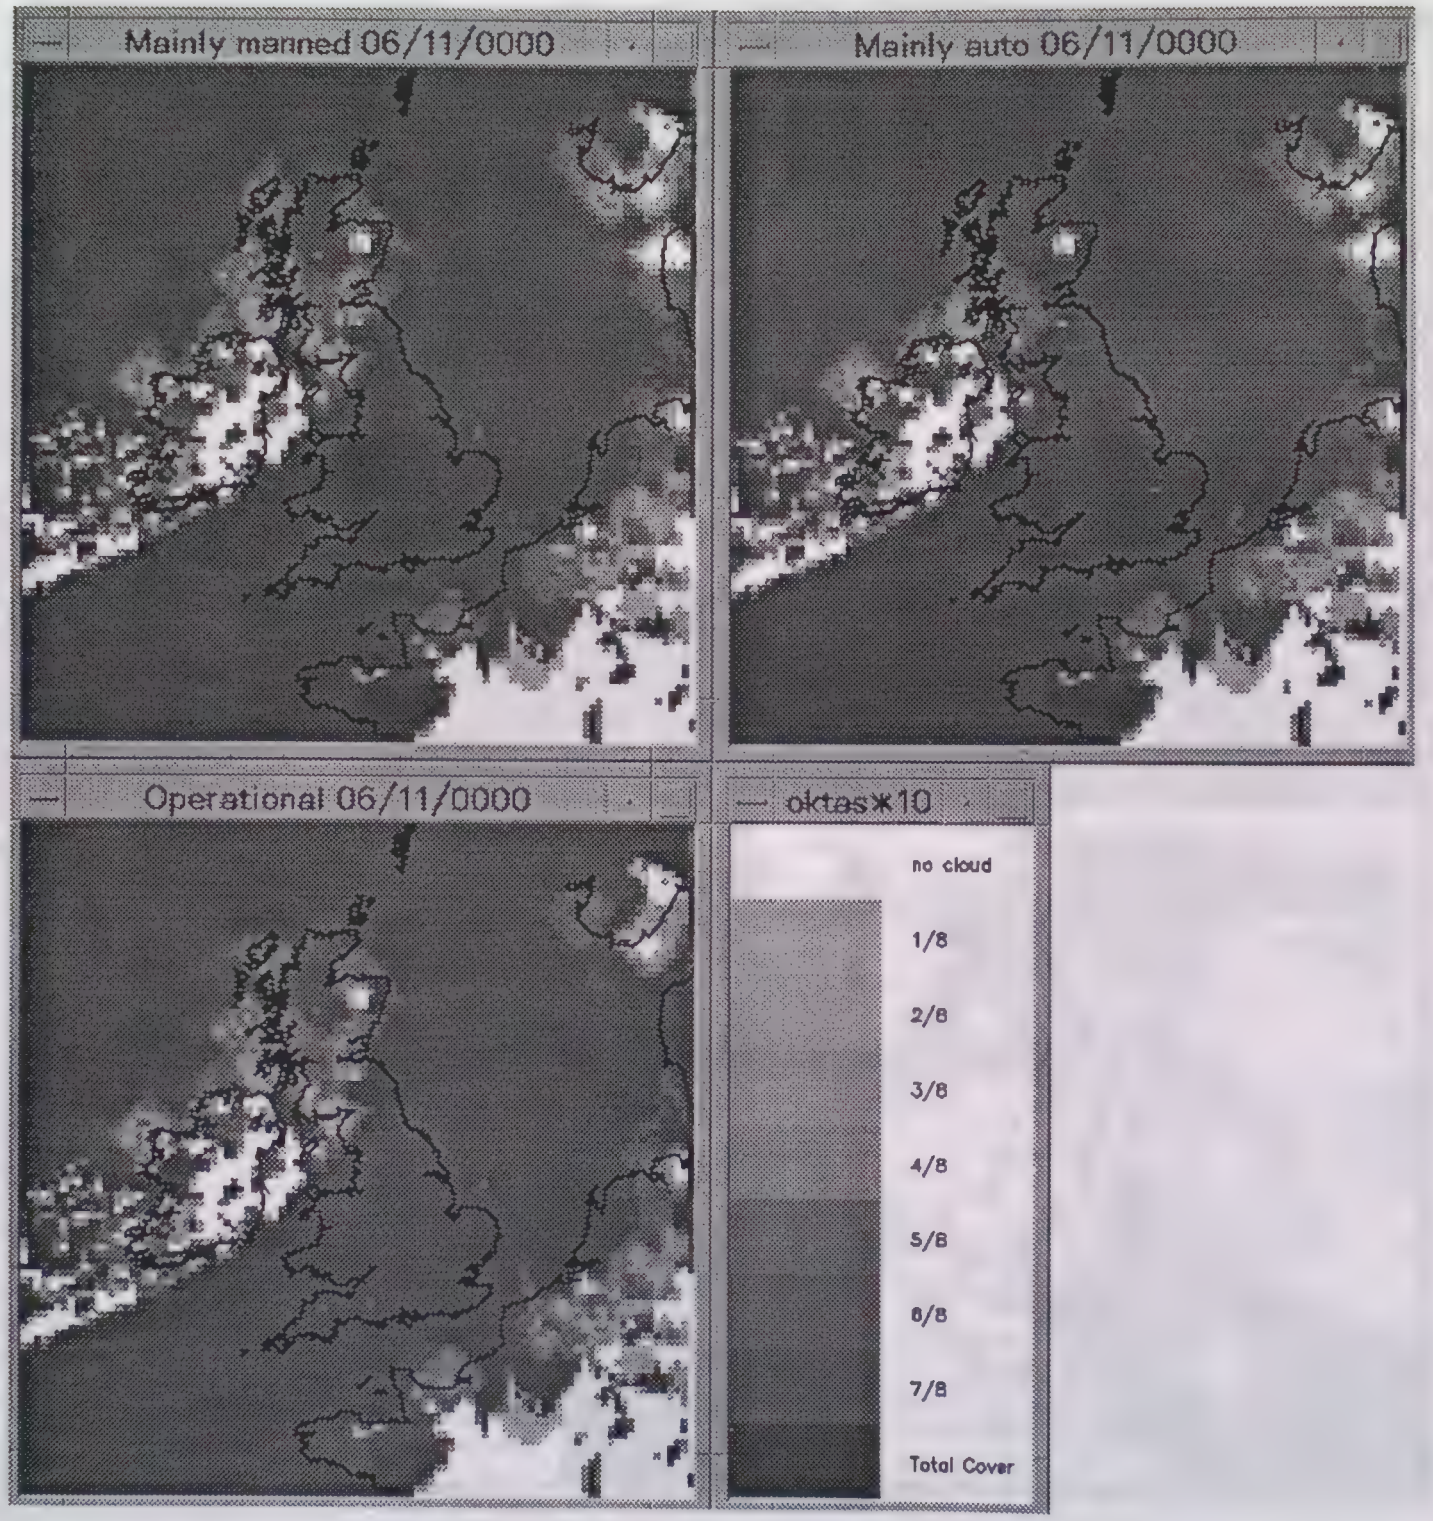

Figure 11. Charts showing amount of total cloud cover in each 15km pixel from the MAN, AUTO-MAN and OP analyses for 0000 UTC 6/11/96. Cloud amounts are shaded according to coverage in oktas.

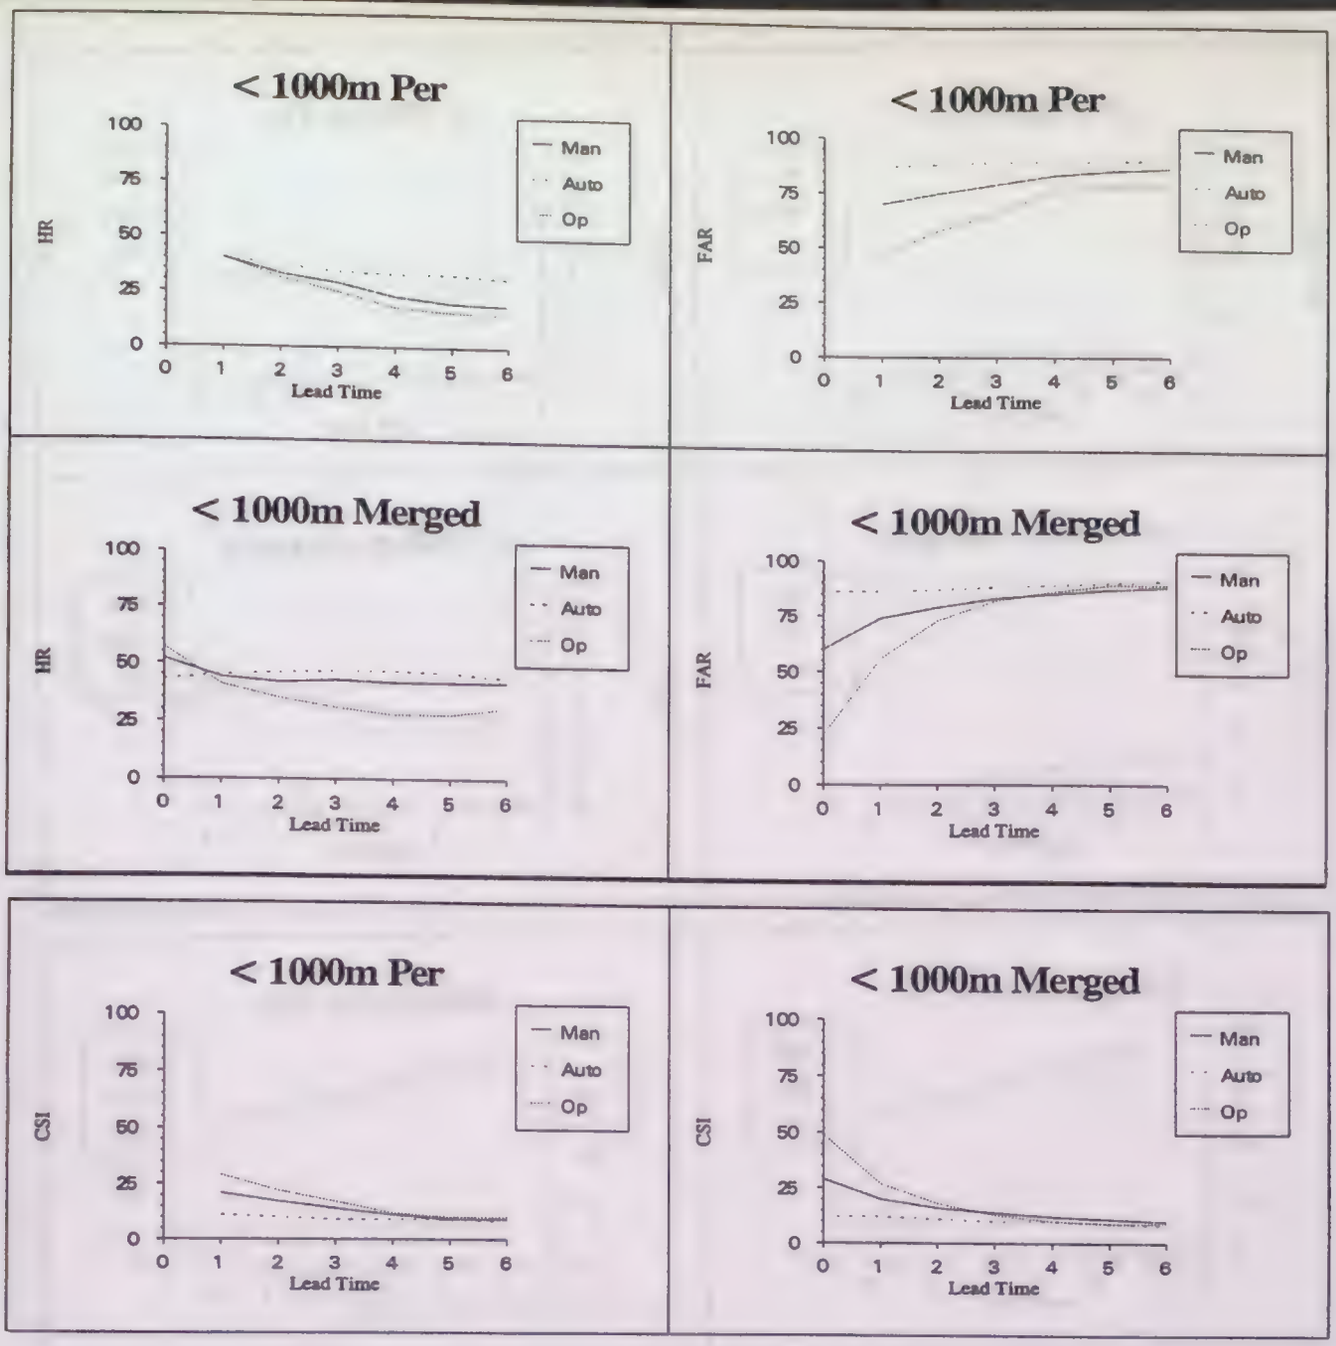

Figure 12. Graphs showing the performance of the operational and test runs for visibility below 1000m in area 1. Hit rate (HR), false alarm rate (FAR) and critical success index (CSI) percentages as a function of lead time are shown for persistence and the Nimrod (merged) analyses and forecasts.

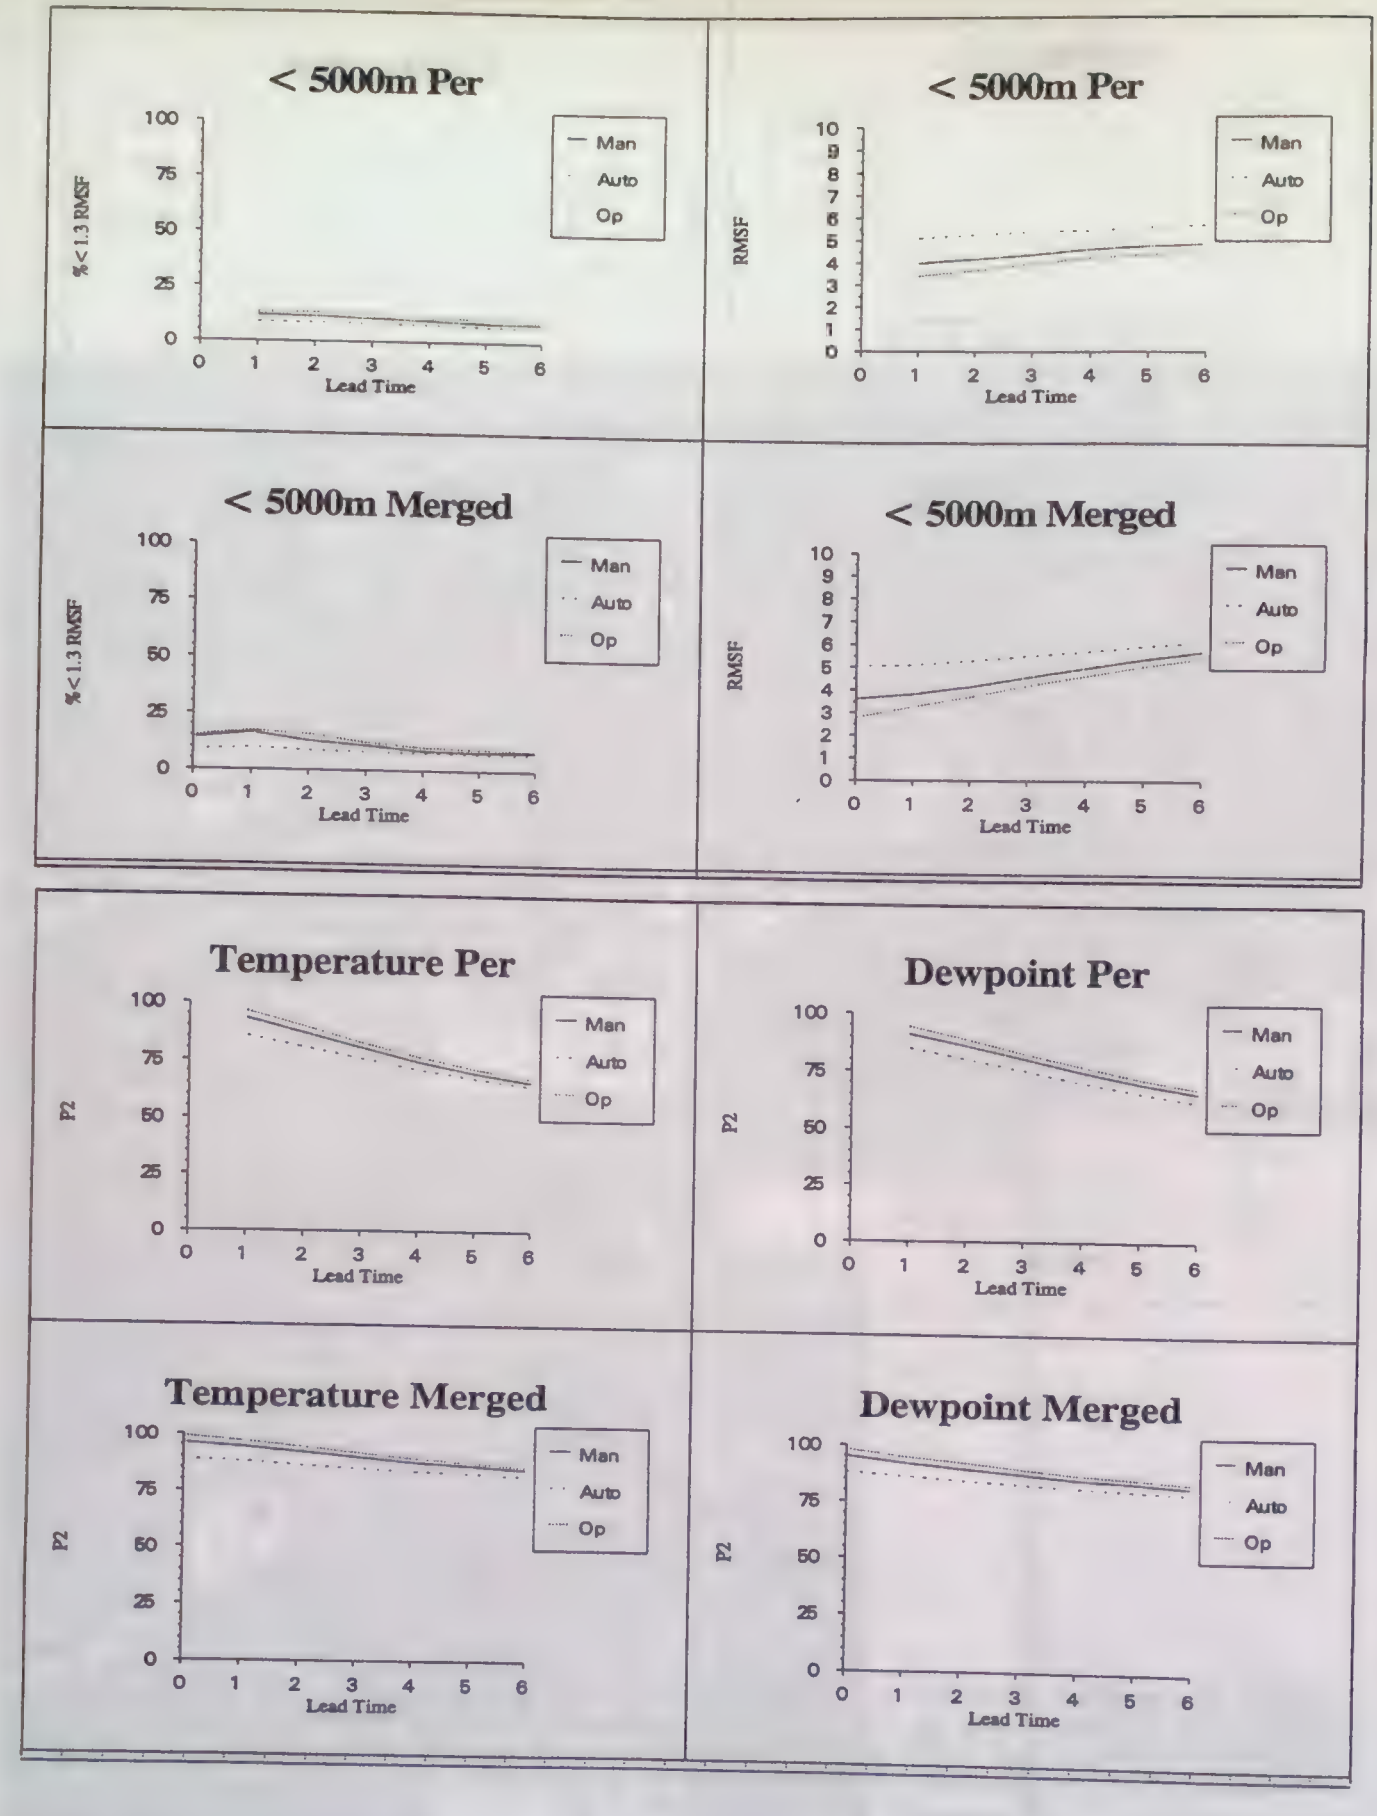

Figure 13. Graphs showing the performance of the operational and test runs for persistence and the Nimrod (merged) analyses and forecasts in area 1. The top graph shows <= 5000m visibility. RMSF and the percentage of occasions when the visibility is within 30% of the observed value (RMSF < 1.3) are plotted as a function of lead time. The bottom graph shows the percentage of occasions when the temperature and dewpoint are within 2 degrees of the observed values.

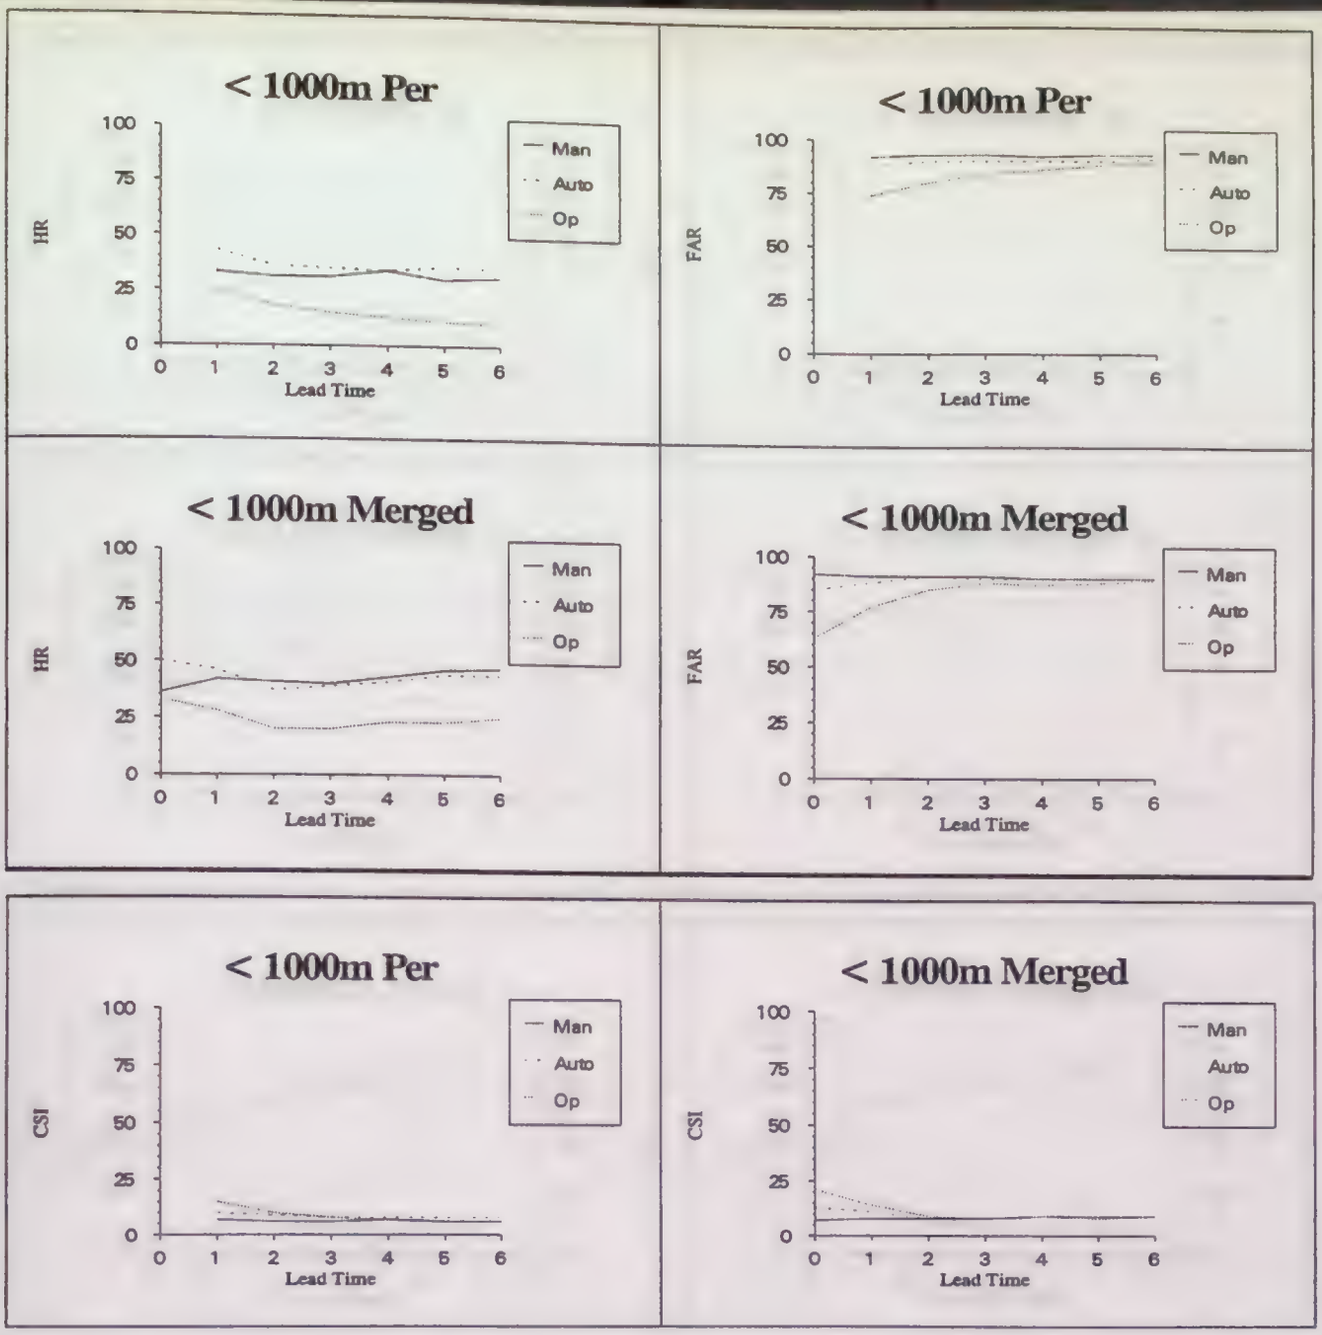

Figure 14. Graphs showing the performance of the operational and test runs for visibility below 1000m in area 2. Hit rate (HR), false alarm rate (FAR) and critical success index (CSI) percentages as a function of lead time are shown for persistence and the Nimrod (merged) analyses and forecasts.

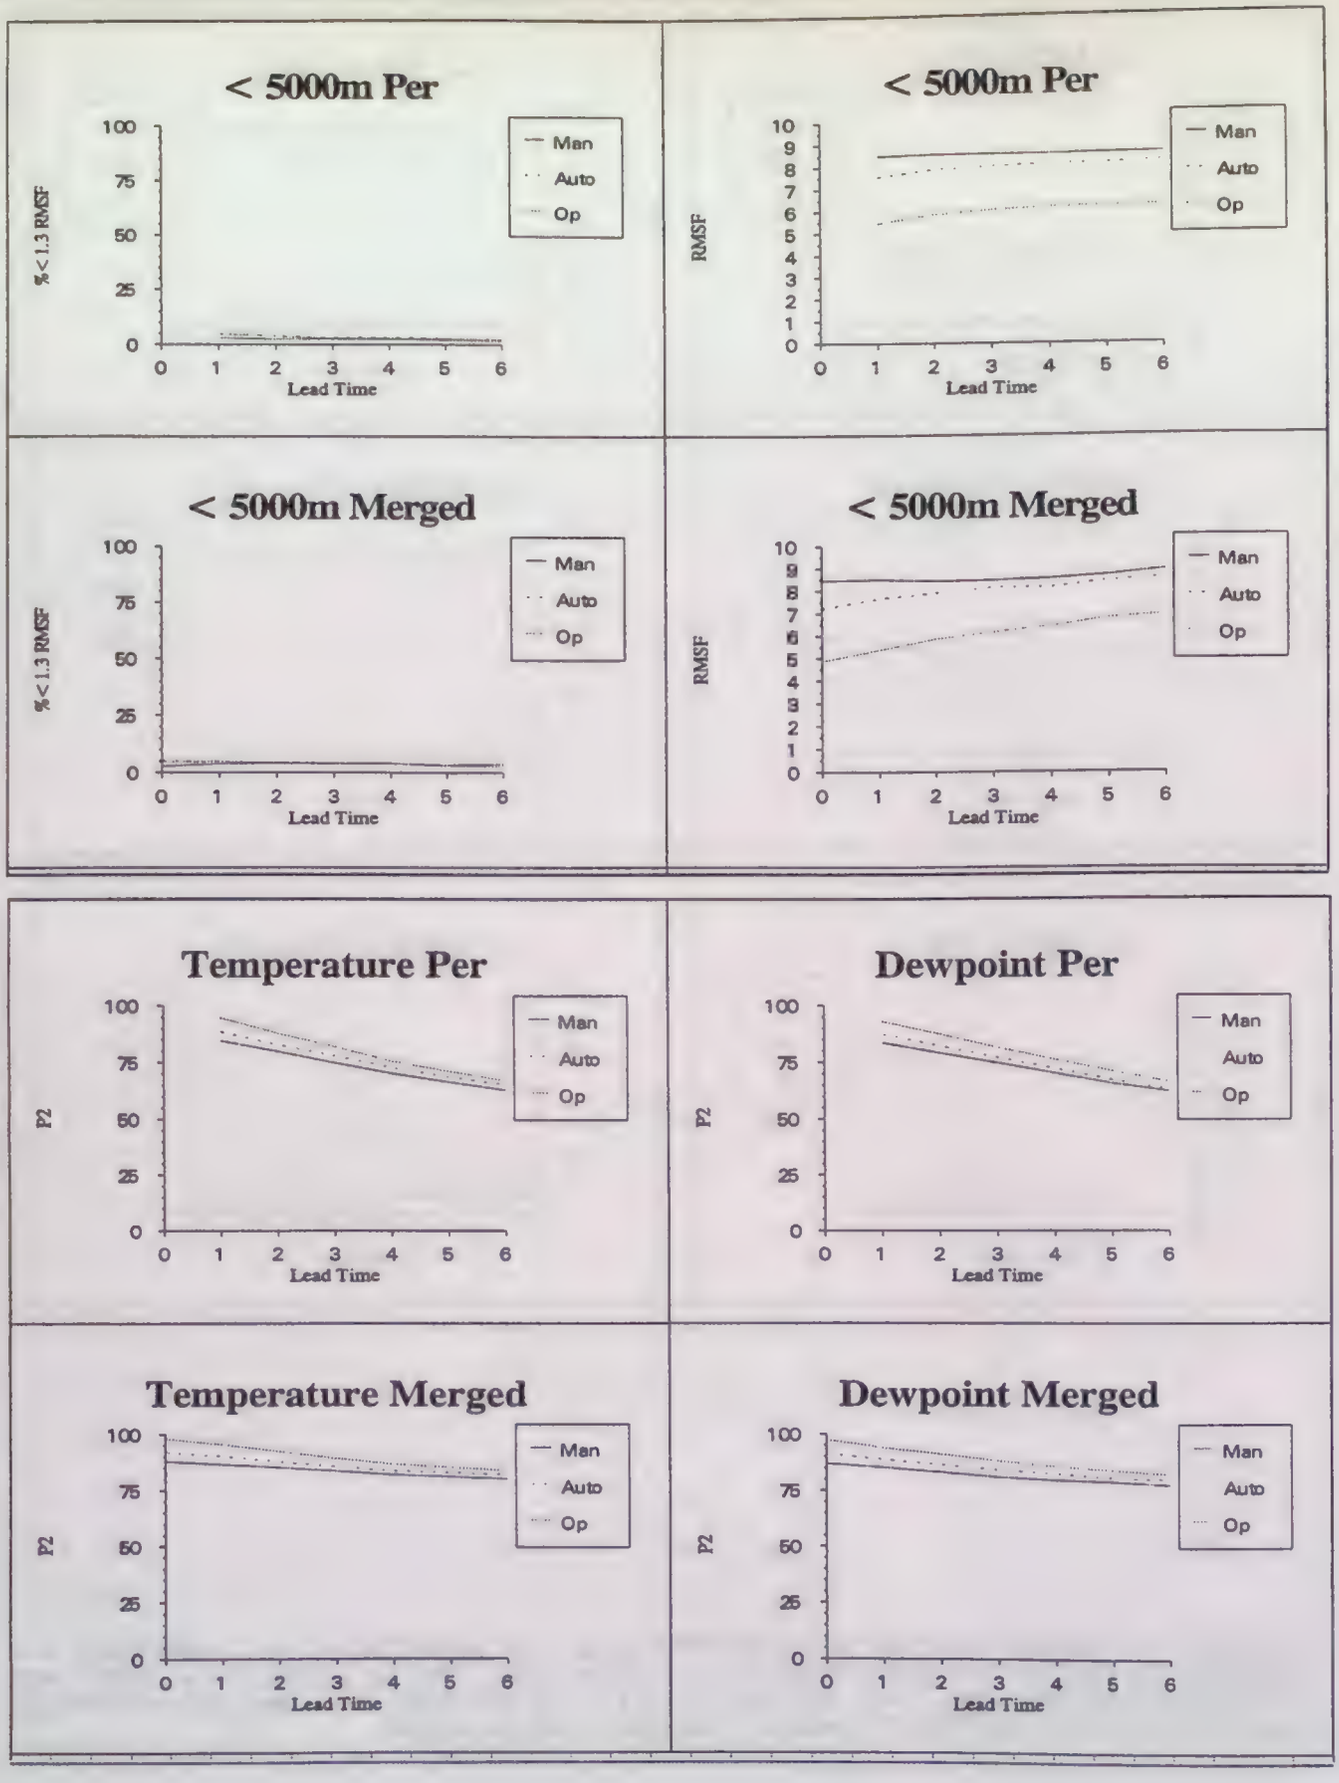

Figure 15. Graphs showing the performance of the operational and test runs for persistence and the Nimrod (merged) analyses and forecasts in area 2. The top graph shows <= 5000m visibility. RMSF and the percentage of occasions when the visibility is within 30% of the observed value (RMSF < 1.3) are plotted as a function of lead time. The bottom graph shows the percentage of occasions when the temperature and dewpoint are within 2 degrees of the observed values.

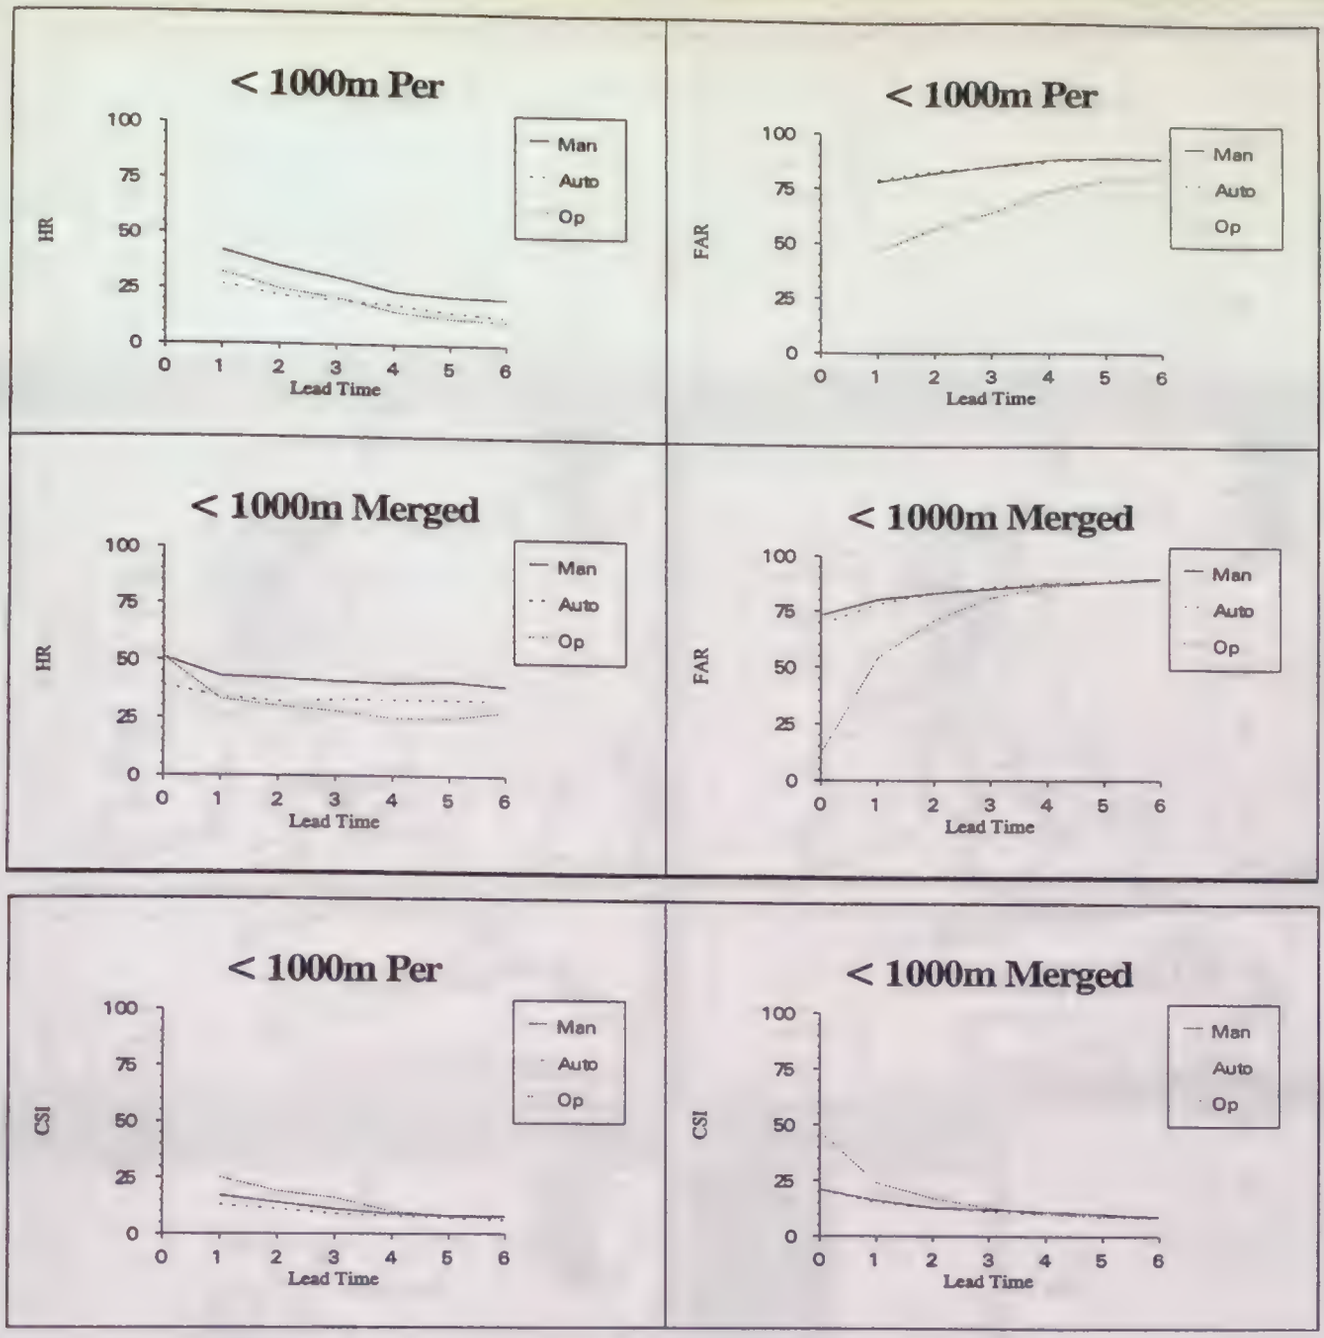

Figure 16. Graphs showing the performance of the operational and test runs for visibility below 1000m in area 3. Hit rate (HR), false alarm rate (FAR) and critical success index (CSI) percentages as a function of lead time are shown for persistence and the Nimrod (merged) analyses and forecasts.

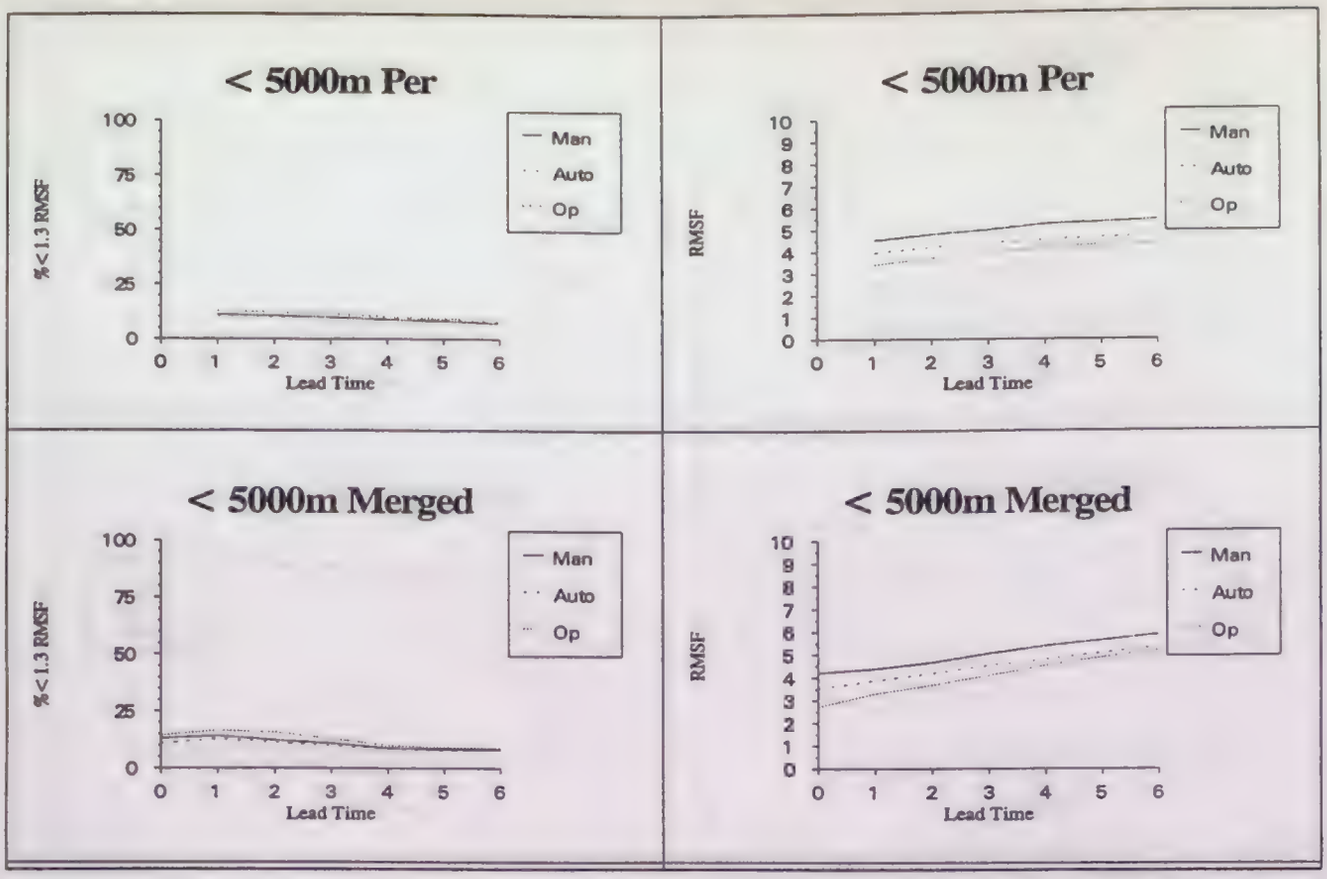

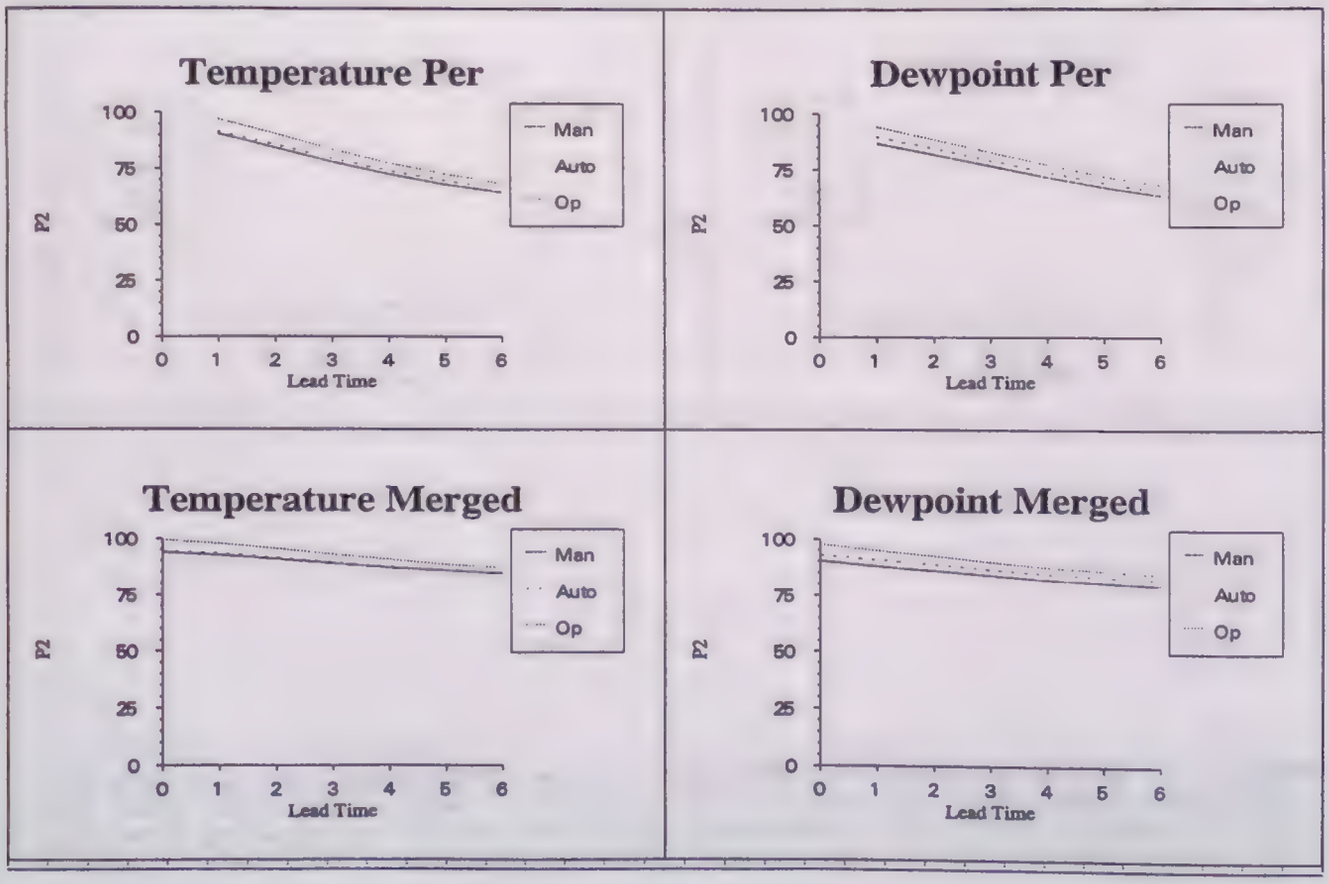

Figure 17. Graphs showing the performance of the operational and test runs for persistence and the Nimrod (merged) analyses and forecasts in area 3. The top graph shows <= 5000m visibility. RMSF, and the percentage of occasions when the visibility is within 30% of the observed value (RMSF < 1.3) are plotted as a function of lead time. The bottom graph shows the percentage of occasions when the temperature and dewpoint are within 2 degrees of the observed values.

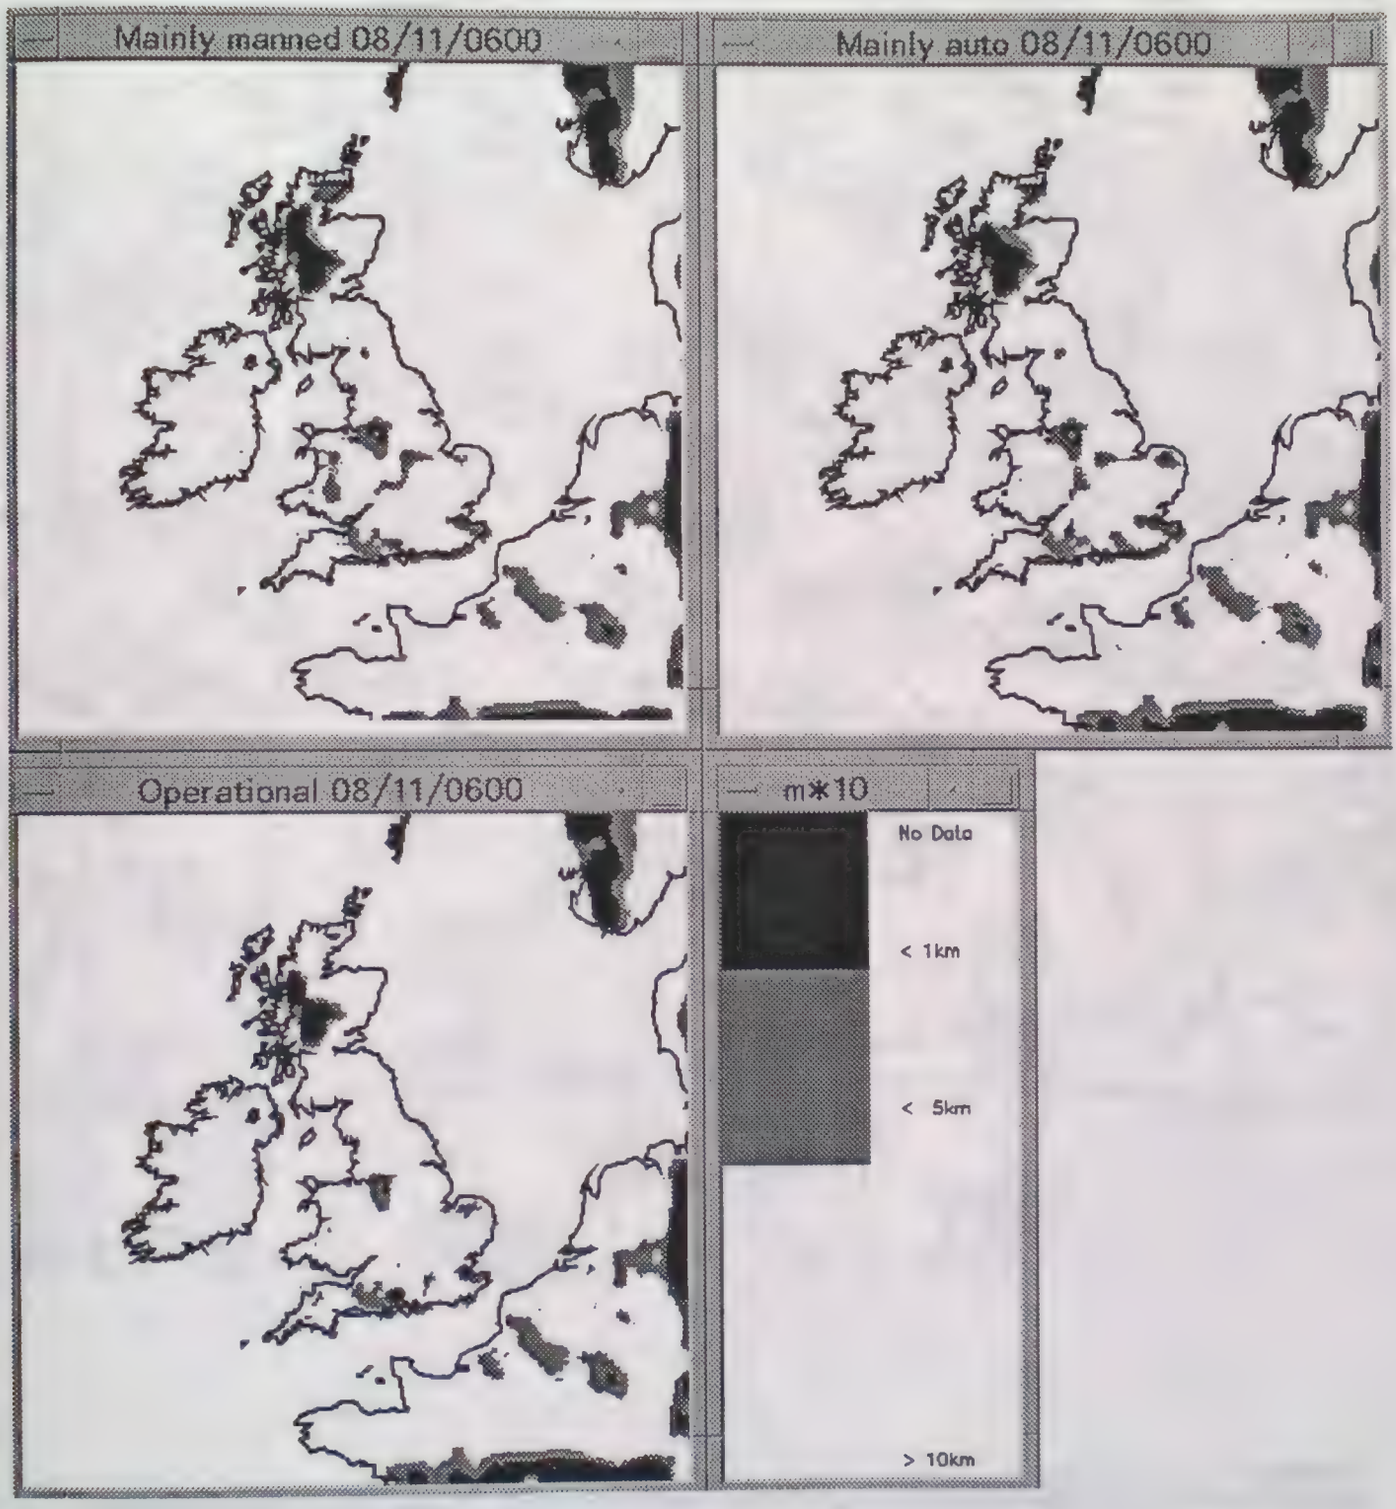

Figure 18. Charts showing Nimrod visibility analyses for thresholds below 1000m (black) and below 5000m (grey) for 0600 UTC 8/11/96 from the MAN, AUTO and OP runs.

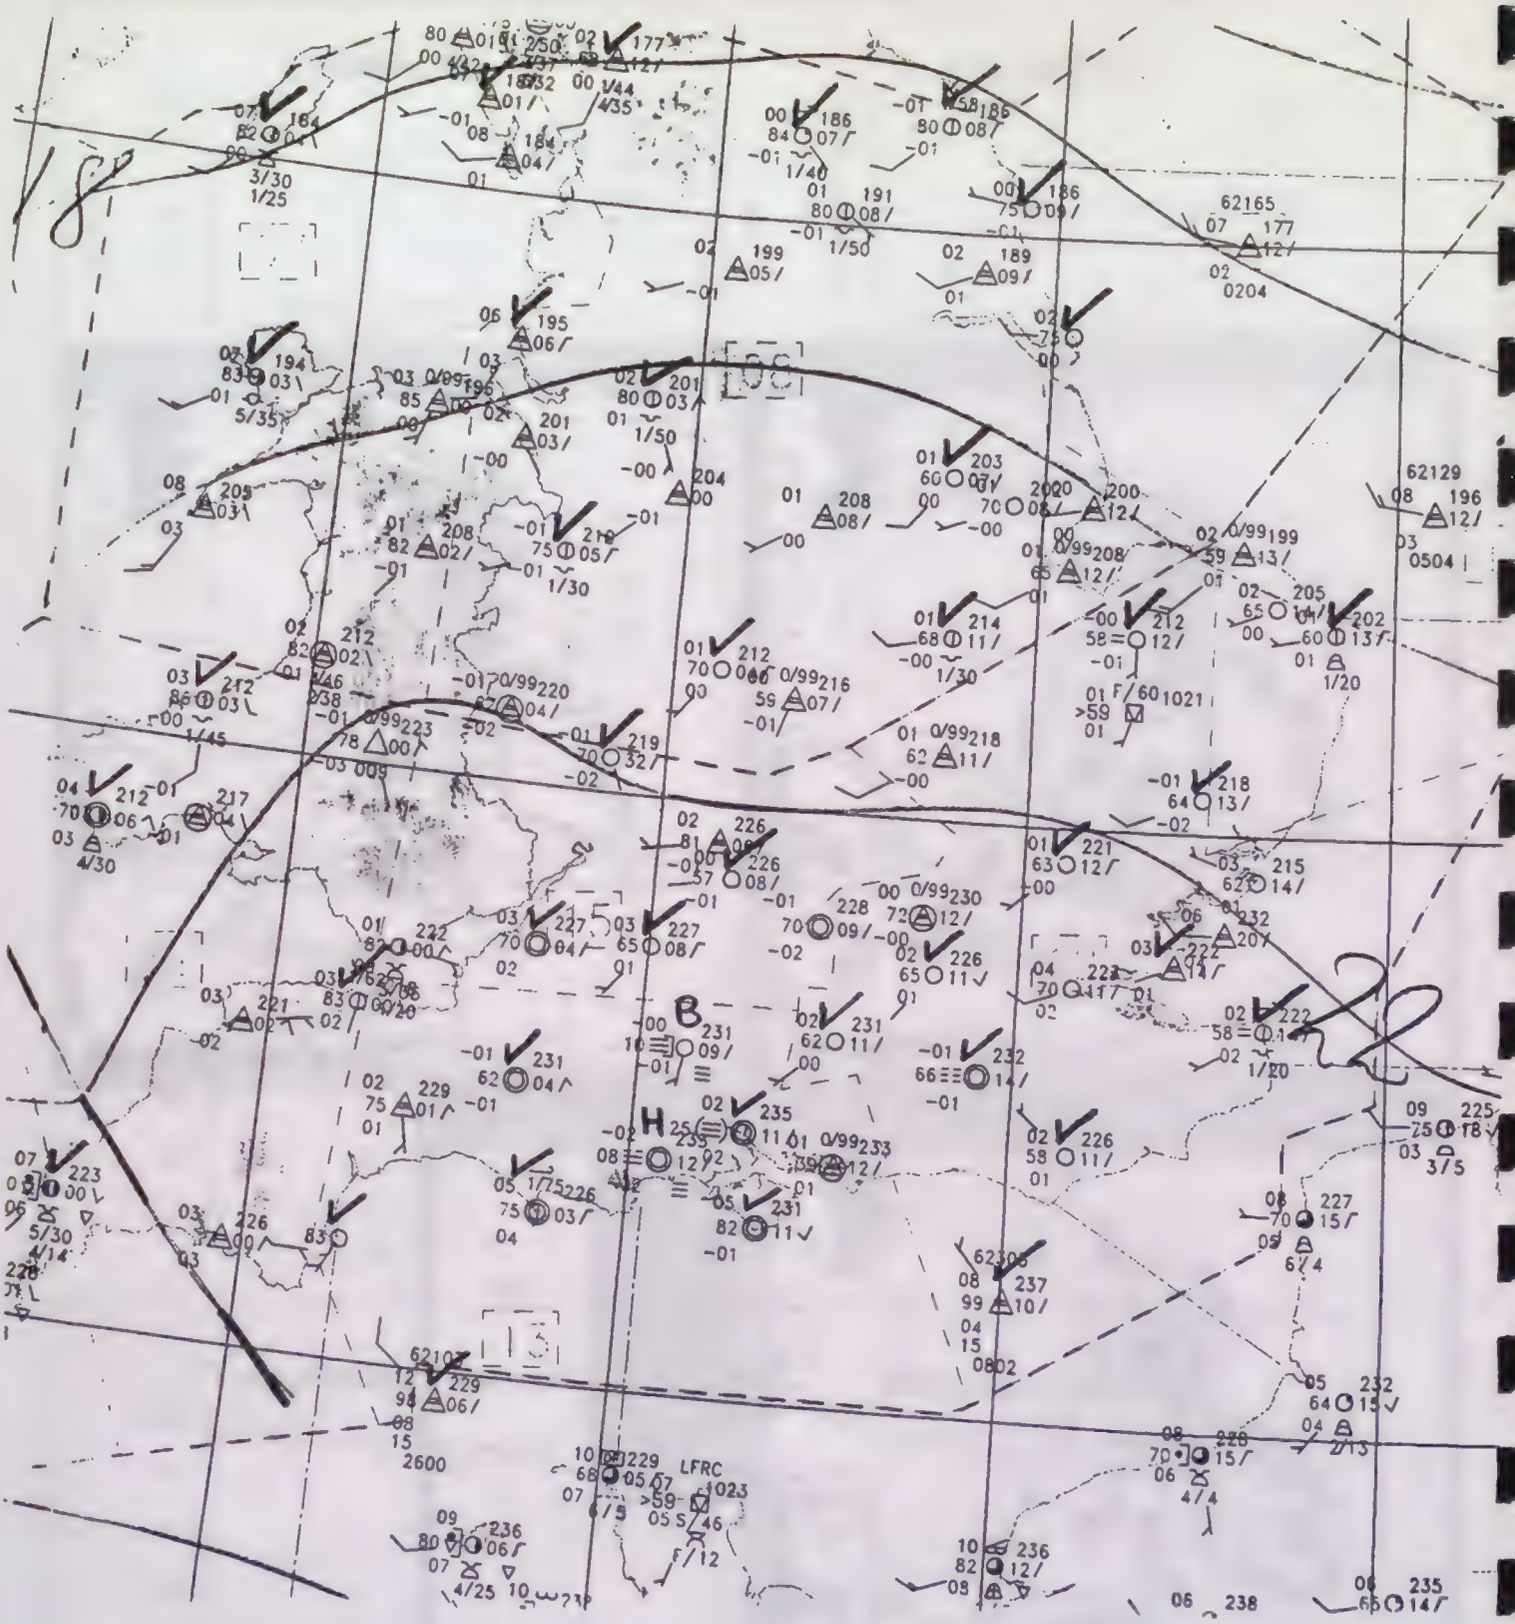

Figure 19. Chart showing plotted observations for 0600 UTC 8/11/96. Those ticked were used in the MAN analysis. 'B' indicates the location of Boscombe Down and 'H' the location of Hurn Airport.

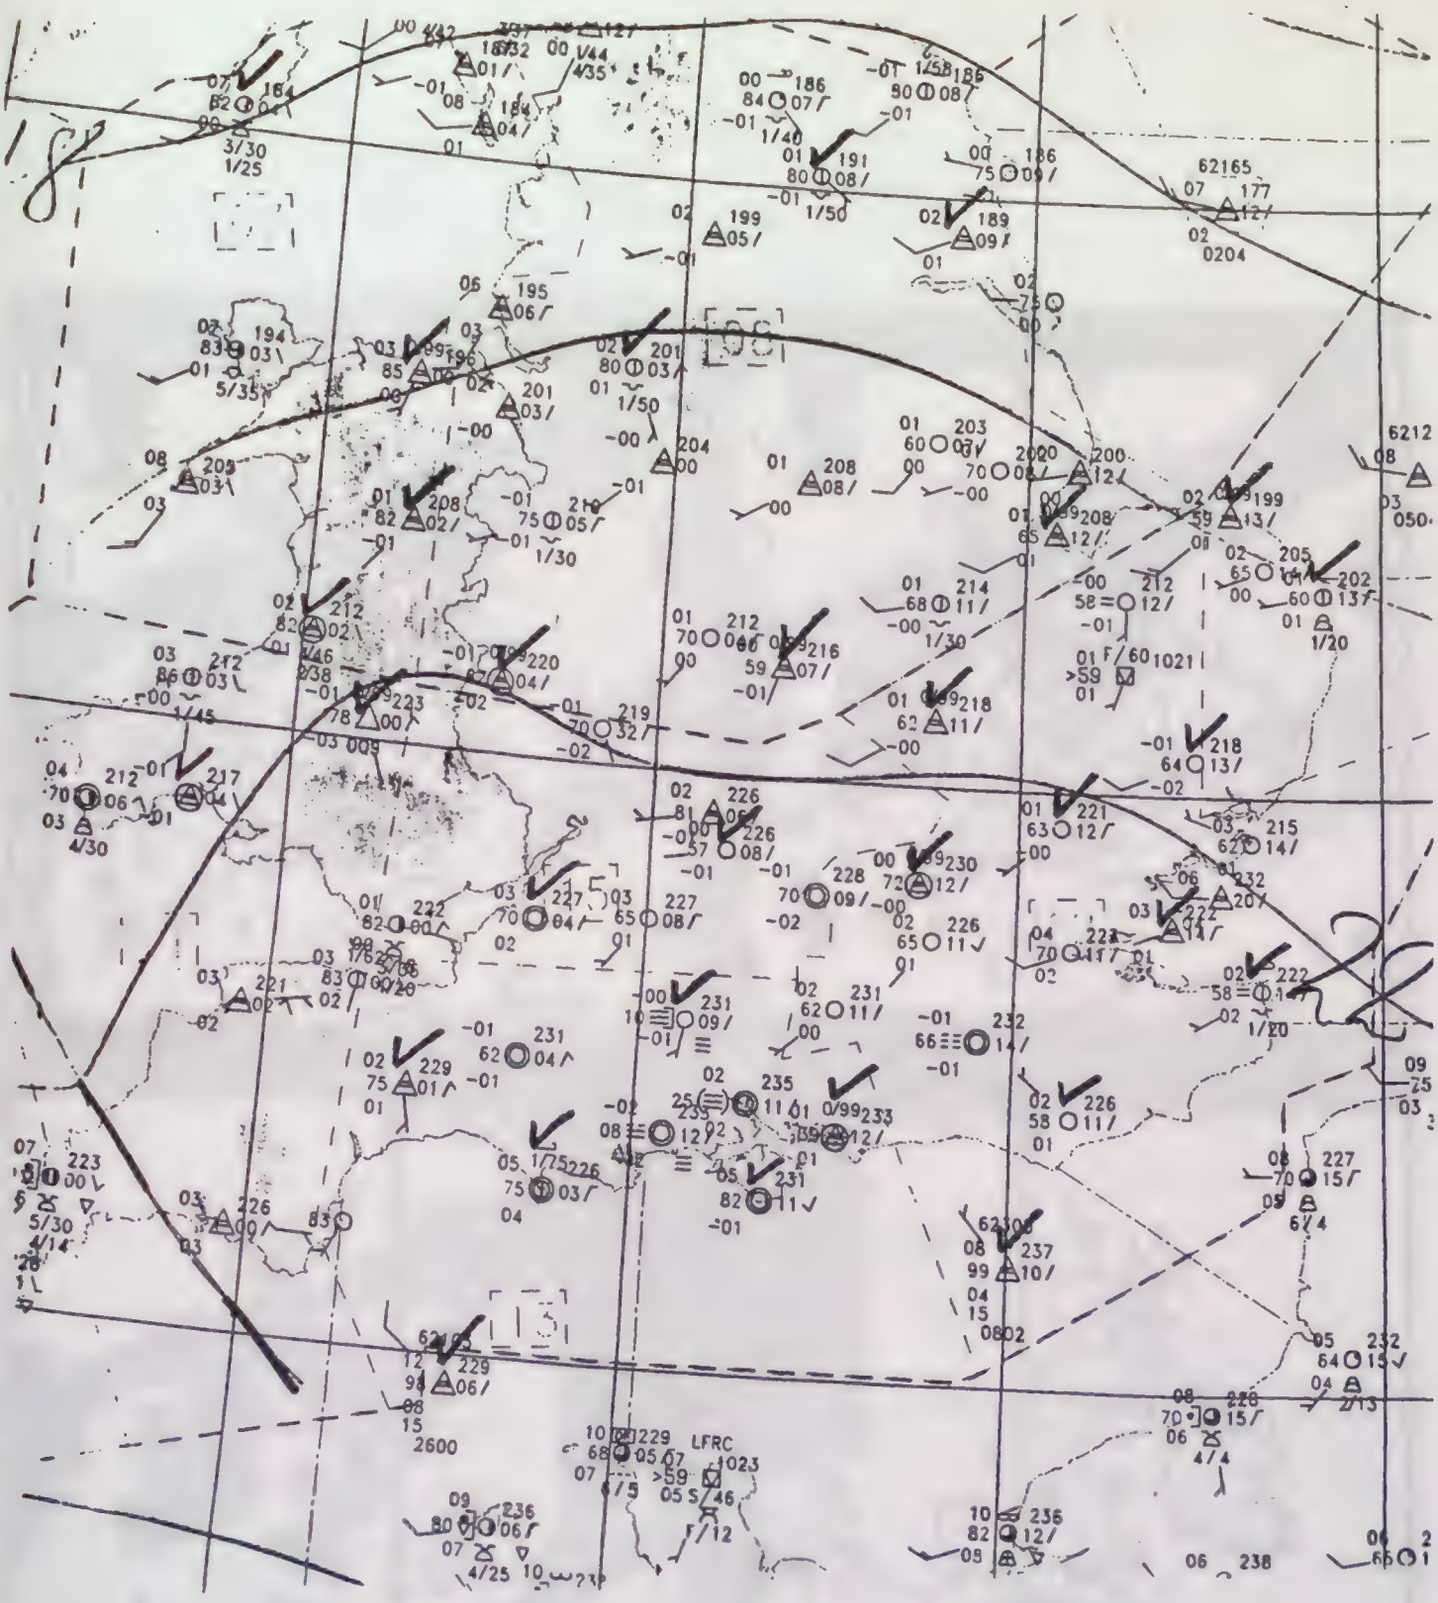

Figure 20. Chart showing plotted observations for 0600 UTC 8/11/96. Those ticked were used in the AUTO analysis.

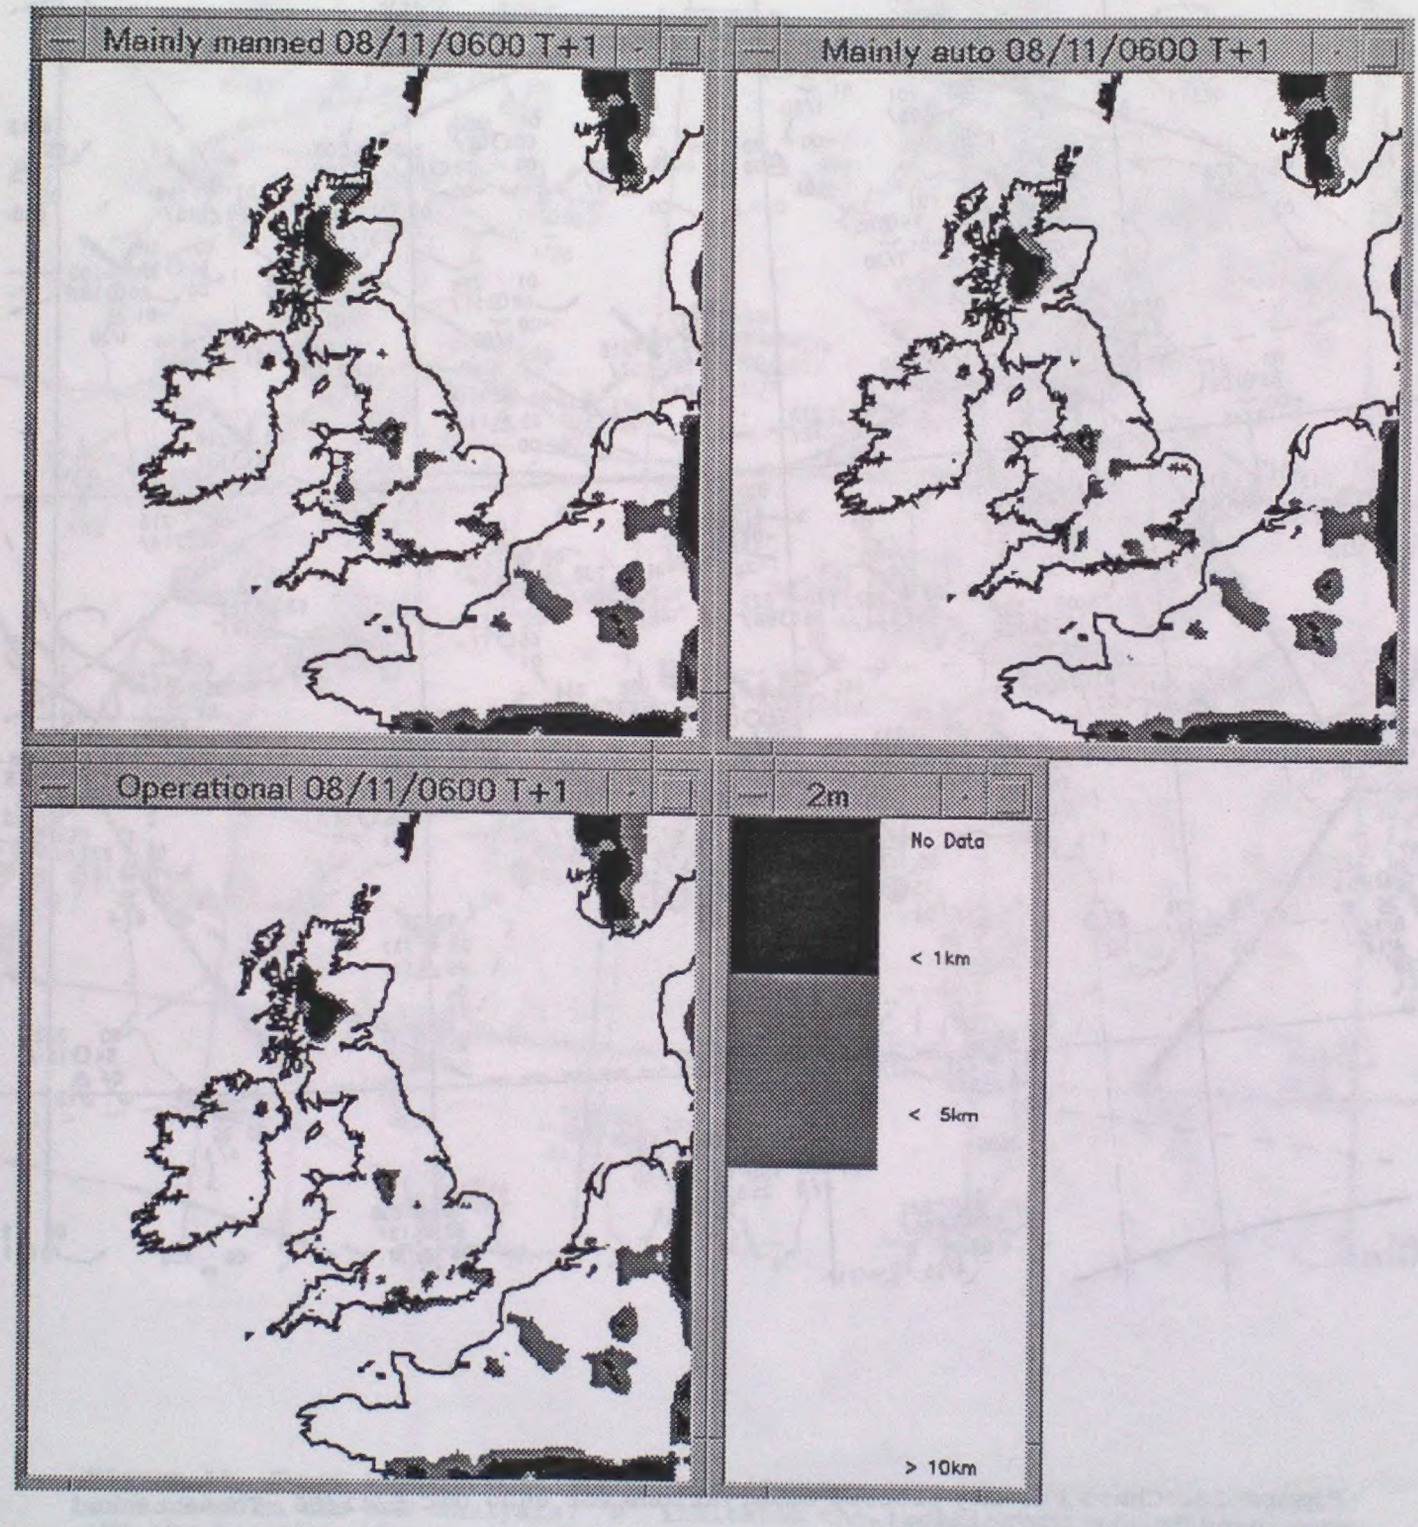

Figure 21. Charts showing Nimrod first guess (T+1) visibility forecasts for 0600 UTC 8/11/96 from the MAN, AUTO and OP runs. Visibility below 1000m is shaded in black and below 5000m in grey.

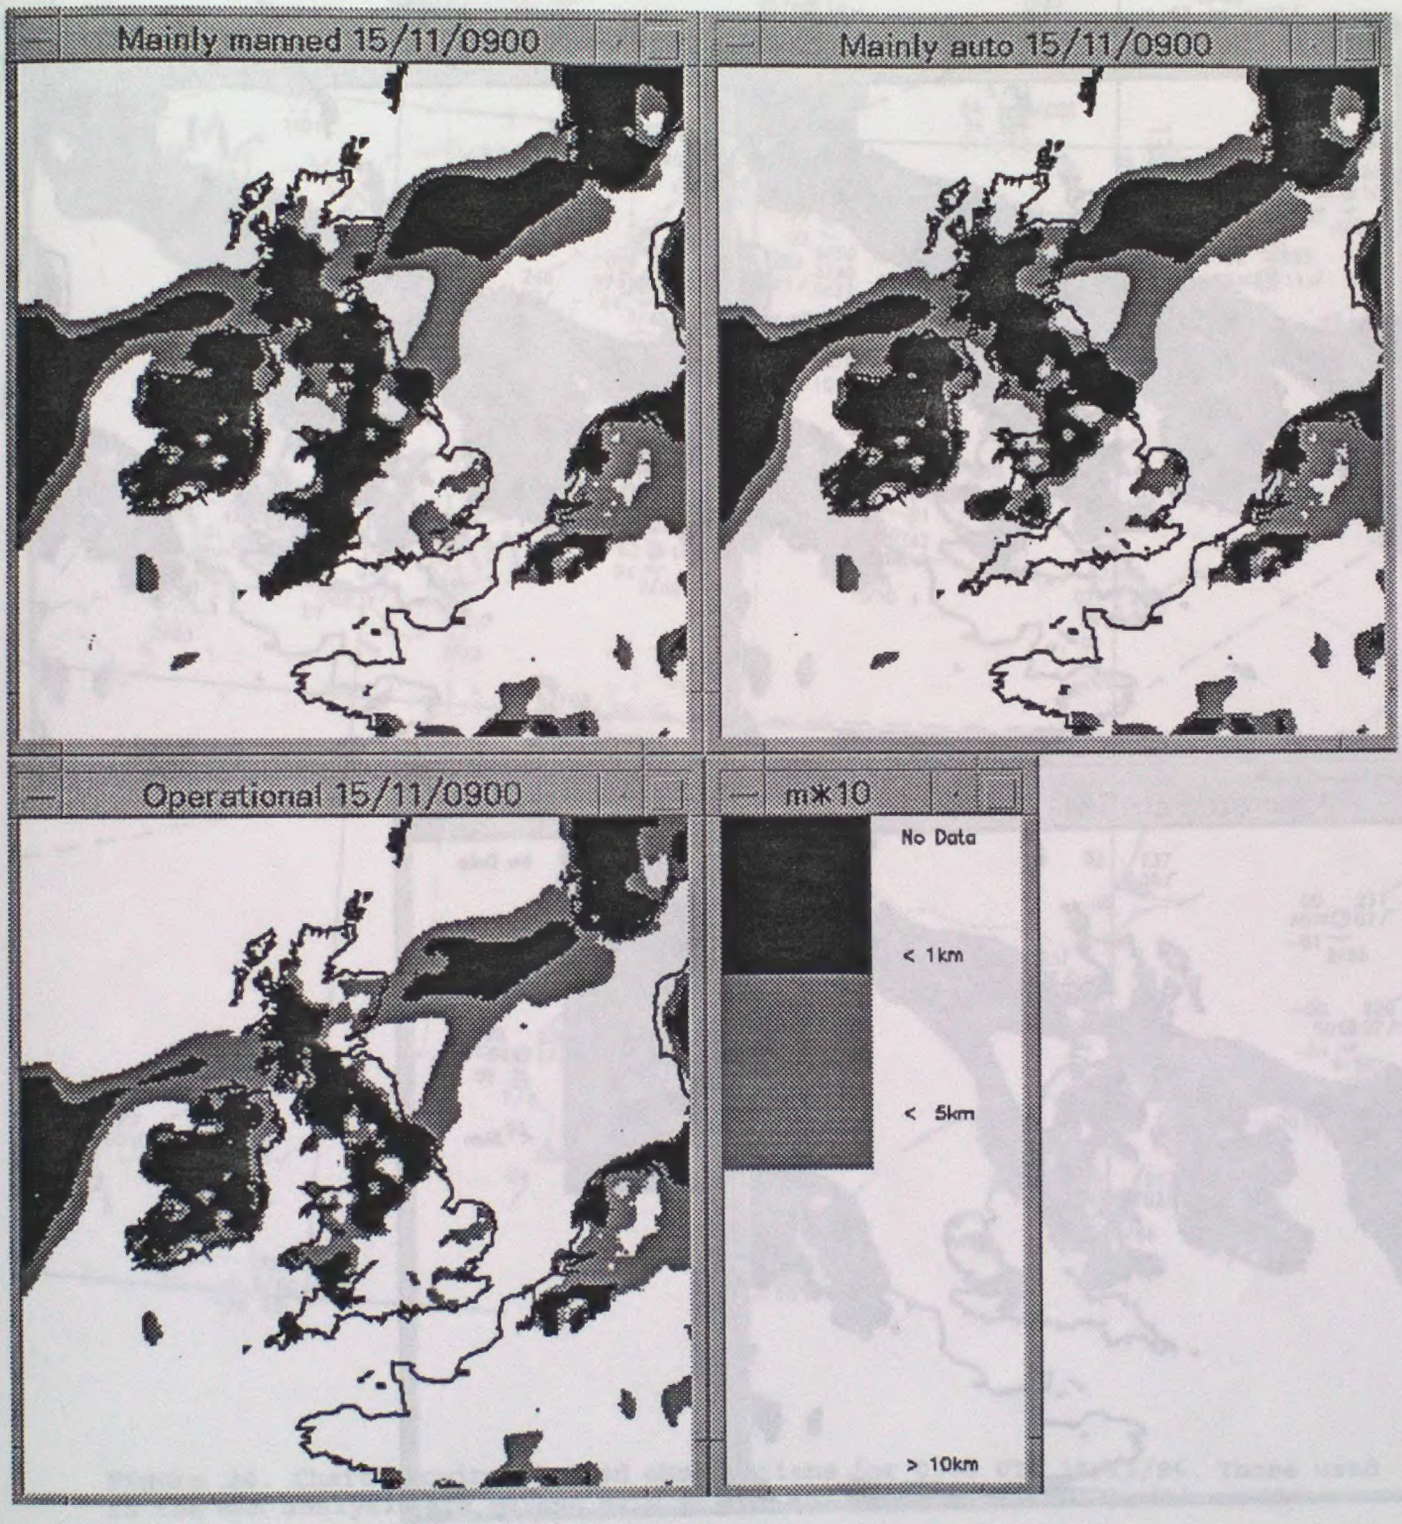

Figure 22. Charts showing Nimrod visibility analyses for thresholds below 1000m (black) and below 5000m (grey) for 0900 UTC 15/11/96 from the MAN, AUTO and OP runs.

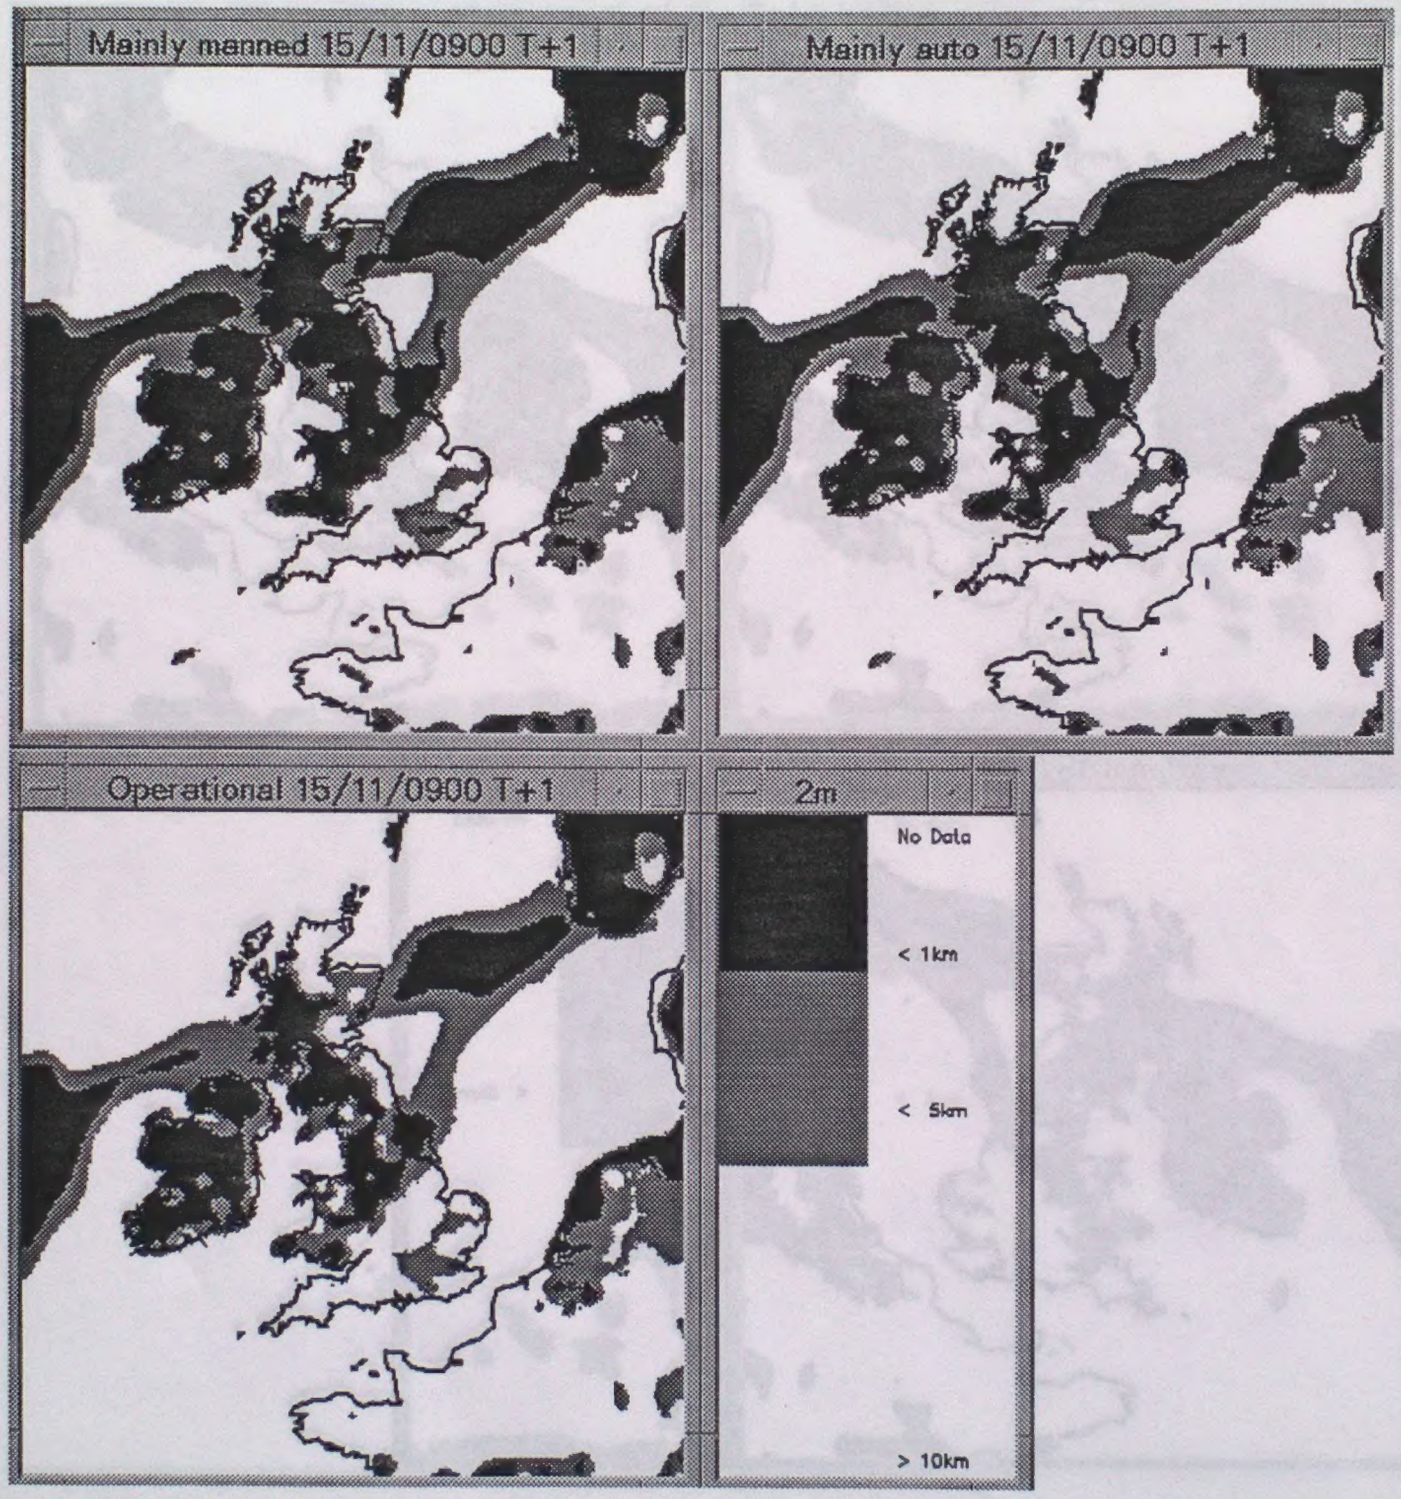

Figure 23. Charts showing Nimrod first guess (T+1) visibility forecasts for 0900 UTC 15/11/96 from the MAN, AUTO and OP runs. Visibility below 1000m is shaded in black and below 5000m in grey.

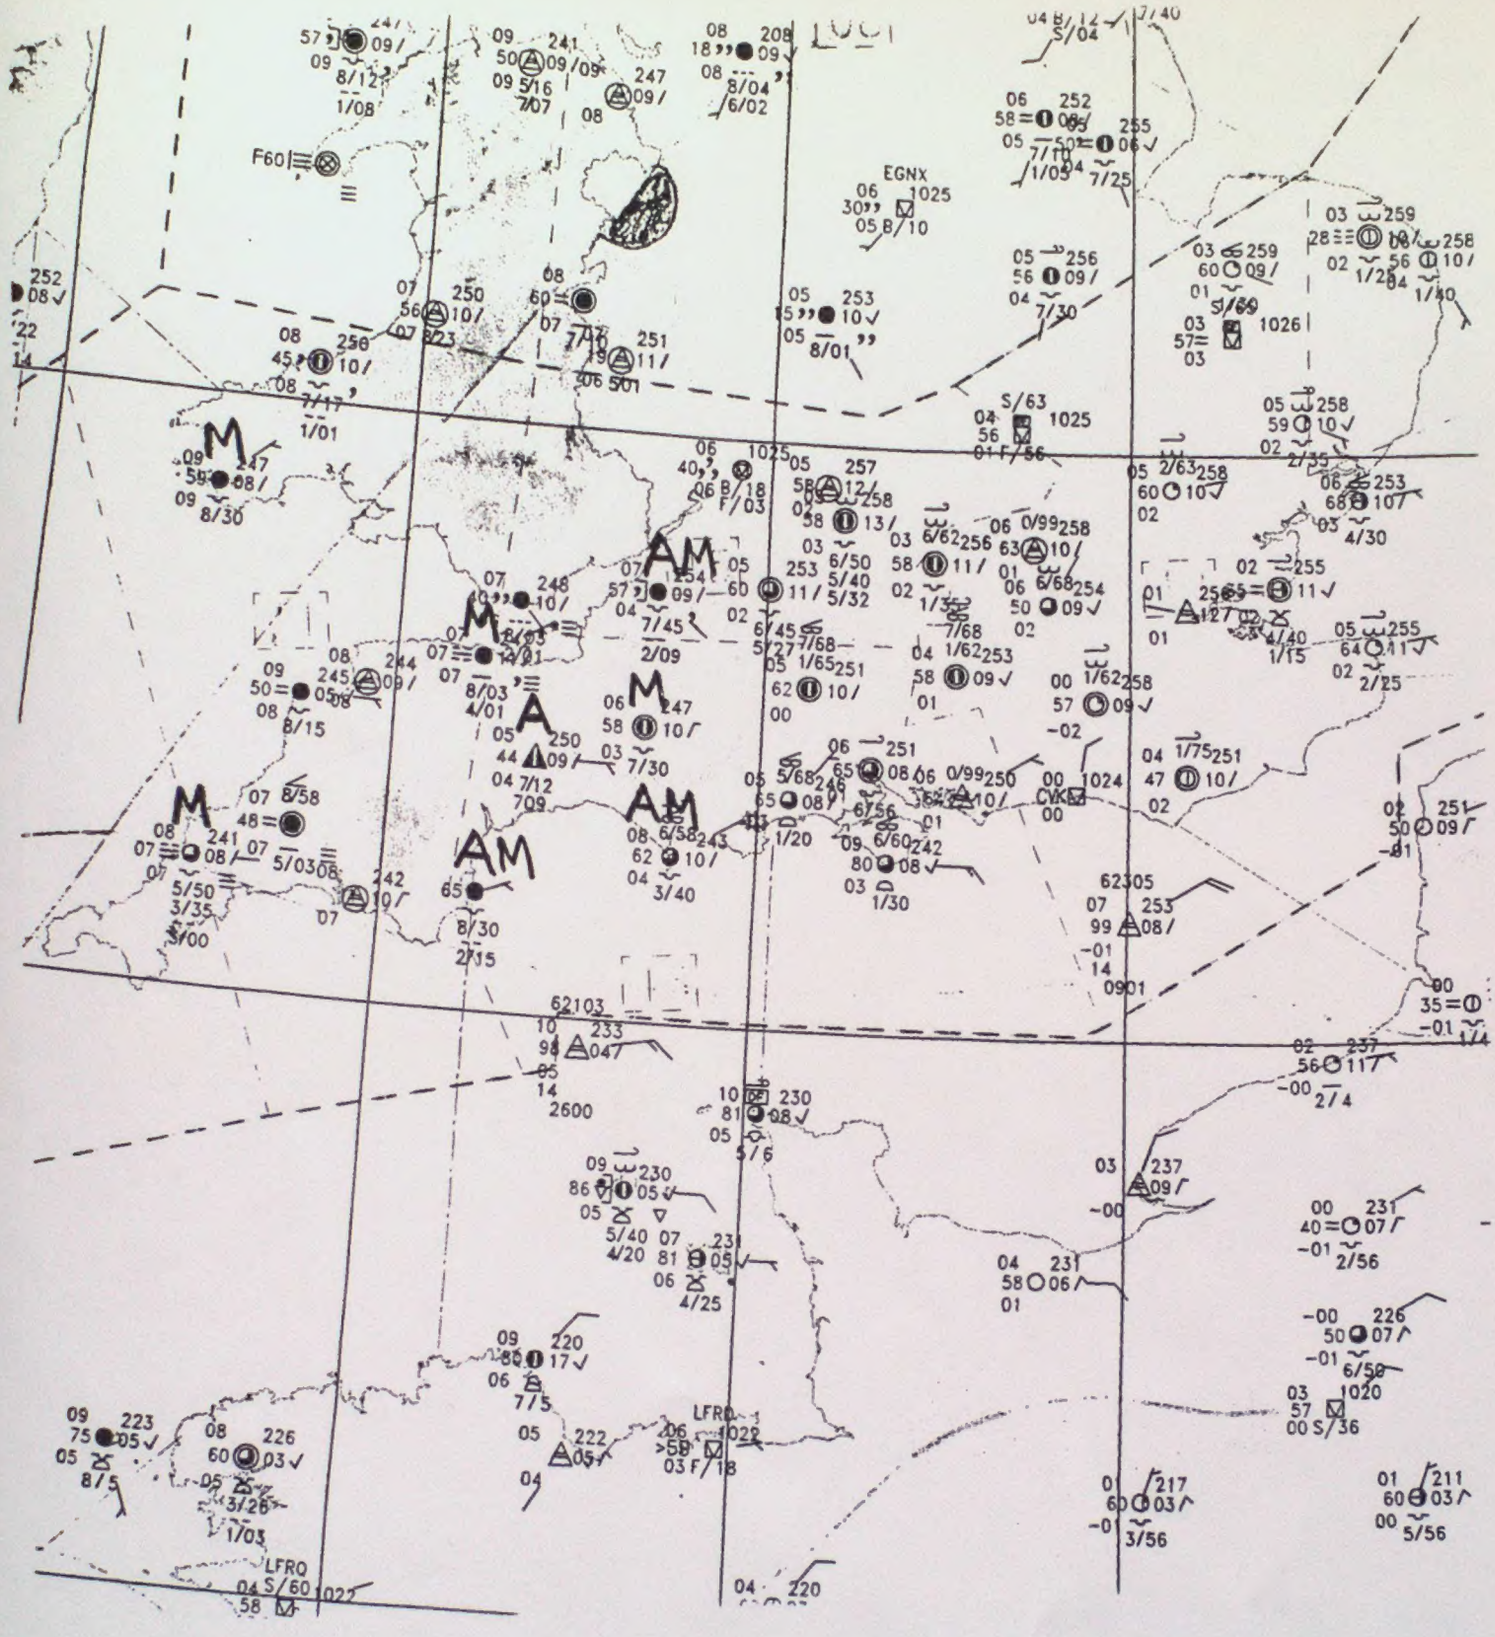

Figure 24. Chart showing plotted observations for 0900 UTC 15/11/96. Those used in the MAN analysis are marked with an 'M' and those in the AUTO with an 'A'.